


CLINICAL INFORMATION AND TECHNOLOGY SERIES

NETWORKS



Michael Bourke, Ph.D.
Stephen Grimes

NETWORKS

Michael Bourke, Ph.D.

College of Business and Economics
Houston Baptist University
7502 Fondren Rd.
Houston, TX 77074-3298

Stephen Grimes

GENTECH
P.O. Box 969
Saratoga Springs, NY 12866

This book is part of the SpaceLabs Medical Clinical Information & Technology Book Series for biomedical and clinical professionals. The series is an educational service of SpaceLabs Medical, a leading provider of patient monitoring and clinical information systems.

© SpaceLabs Medical, Inc., 1995

All rights reserved.

No part of this book may be reproduced by any means or transmitted, or translated into a machine language without the written permission of the publisher.

All brands and product names are trademarks of their respective owners.

Published by SpaceLabs Medical, Inc., Redmond, Washington, U.S.A.

Printed in the United States

ISBN 0-9627449-9-9

TABLE OF CONTENTS

	Page		Page
INTRODUCTION	1	2.7.2 De Facto Standards - TCP/IP	21
1.0 NETWORKING	2	2.7.3 Vendor Conventions - SPX	21
1.1 Local Area Networks	2	2.8 OSI Layer 5 - Session	22
1.2 Wide Area Networks	2	2.8.1 Middleware	22
1.3 Internetworks	3	2.9 OSI Layer 6 - Presentation	23
1.4 PC LANs	3	2.10 OSI Layer 7 - Application	23
1.5 Recent Trends	4	2.10.1 Naming	25
1.5.1 Network Integration	4	2.10.2 Electronic Mail	25
1.5.2 Wireless LANs	6	2.10.3 Complex Document Exchange	26
1.5.3 Higher Throughput	6	2.10.4 Remote Procedure Call	27
1.5.4 Network Management Tools	7	2.10.5 File Exchange	27
1.6 Network Operating Systems	7	2.10.6 Remote Database Access	28
2.0 OPEN SYSTEMS INTERCONNECTION	7	2.10.7 Electronic Data Interchange	28
2.1 Standards Organizations	8	3.0 NETWORK INTERFACES	30
2.1.1 International Organizations for Standardization and International Electrotechnical Committee	8	3.1 Cabling	31
2.1.2 International Association of Electrical and Electronic Engineers	9	3.1.1 Coaxial Cable	31
2.1.3 International Telephone and Telegraph Consultative Committee	9	3.1.2 Twisted-Pair Cable	31
2.1.4 Open System Foundation	9	3.1.2.1 Unshielded Twisted-Pair	32
2.1.5 Corporation for Open Systems	9	3.1.2.2 Shielded Twisted-Pair	32
2.2 Open Systems Interconnection Model	10	3.1.3 Fiber-Optic Cable	32
2.3 The OSI Model and Biomedical Devices	12	3.1.4 Wireless Connections	32
2.4 OSI Layer 1 - Physical	14	3.1.4.1 Infrared Transceivers	33
2.4.1 Topology	15	3.1.4.2 Radio Signals	33
2.5 OSI Layer 2 - Data Link	15	3.1.4.3 Microwave Signals	33
2.6 OSI Layer 3 - Network	16	3.2 Topologies	33
2.6.1 International Standards	17	3.2.1 Bus Topology	33
2.6.2 De Facto Standards	18	3.2.2 Star	34
2.6.3 Vendor Conventions	18	3.2.3 Ring	34
2.6.3.1 IBM Systems Network Architecture	18	3.3 Wiring Centers	36
2.6.3.2 DEC Local Area Transport	19	3.3.1 Hubs	37
2.6.3.3 Novell Internet Packet Exchange/Sequenced Packet Exchange	19	3.3.2 Concentrators	37
2.6.4 Tunneling	19	3.3.3 Multistation Access Units	37
2.7 OSI Layer 4 - Transport	21	3.3.4 Repeaters	37
2.7.1 International Standards	21	3.4 LAN Standards/Architectures	37
		3.4.1 IEEE 802.3 - Ethernet	37
		3.4.1.1 10Base5 Standard	38
		3.4.1.2 10Base2 Standard	38
		3.4.1.3 10BaseT Standard	40
		3.4.2 Token Ring Standard	41
		3.4.3 ARCnet Architecture	42
		3.5 Types of Networks	44
		3.5.1 Host-Terminal Networks	44
		3.5.2 Client/Server Networks	44
		3.5.3 Peer-to-Peer Networks	46
		3.6 Network Operating Systems	46

	Page		Page
3.6.1 Novell NetWare	46	6.3 Components to be Managed	62
3.6.2 Microsoft LAN Manager and Windows NT Server	47	6.3.1 Device Management	62
3.6.3 Banyan Systems VINES	47	6.3.2 Simple Network Management Protocol	62
3.6.4 Performance Technology POWERserve	48	6.3.3 OSI CMIP	63
3.6.5 Digital Equipment DECnet and Pathworks	48	6.4 Vendor Solutions	65
3.6.6 DOS-based LANs	48	6.4.1 Netview	65
3.6.7 UNIX	49	6.4.2 Open View	65
4.0 NETWORK APPLICATIONS	49	6.4.3 SUN Net Manager	65
4.1 Resource Sharing	49	6.4.4 Other Managers	65
4.2 Electronic Mail	50	6.5 Cable Management	66
4.3 Group Scheduling (Resource Management)	50	6.5.1 Cable Scanners	66
4.4 Wide Area Networking	50	6.5.2 Protocol Analyzers	66
4.4.1 Scheduled or On-Demand Connections	50	6.5.3 Cable Management System	66
4.4.2 Continuous, Dedicated Connections	51	6.6 Application Management	66
4.4.2.1 T-1 Lines	51	6.6.1 Application Metering Software	66
4.4.2.2 Integrated Services Digital Network	51	6.6.2 Security and Viruses	67
4.4.2.3 Metropolitan Area Networks ...	51	6.7 Workstation Inventory Management	67
4.4.2.4 Fiber Distributed Data Interface	52	7.0 EXTENDED NETWORKS	68
5.0 NETWORK HARDWARE	52	7.1 Technology for Extended Networks	68
5.1 Servers	52	7.2 Community Healthcare Information Network	69
5.1.1 File Server	52	7.3 Public Access Systems	70
5.1.2 Print Server	53	8.0 HEALTHCARE STANDARDS	70
5.1.3 Communications Server	54	8.1 HL7	71
5.1.4 Application Server	55	8.2 MEDIX	72
5.2 Workstations	55	8.3 ASTM	74
5.3 Backup Systems	56	8.4 ACR/NEMA	75
5.4 Network Security	56	8.5 MIB	75
5.5 Fault-Tolerant Systems	57	8.6 Vendor Implementations	75
5.5.1 Uninterruptible Power Supplies	57	9.0 CASE STUDIES	76
5.5.2 Redundant Systems	58	9.1 M. D. Anderson Cancer Center	76
5.5.2.1 RAID 1	58	9.1.1 Overview of Networking	76
5.5.2.2 RAID 2	58	9.1.2 Networks and Applications	78
5.5.2.3 RAID 3	59	9.1.3 Budget and Plans	84
5.5.2.4 RAID 4	59	9.2 Sisters of Charity	85
5.5.2.5 RAID 5	59	9.2.1 Corporate Headquarters	87
5.6 Bridges	59	9.2.2 St. Mary's Hospital	90
5.7 Routers	60	10.0 ABBREVIATIONS	93
5.8 Gateways	60	11.0 BIBLIOGRAPHY	95
6.0 NETWORK MANAGEMENT	60	12.0 GLOSSARY	96
6.1 Single Vendor vs. Multivendor Environments	61	INDEX	99
6.2 Consultants	61		

INTRODUCTION

The objective of this publication is to provide the biomedical engineering and hospital information systems staff with a general overview and basic reference regarding computer networking and its relevance to the hospital environment. The intent is not to provide an in-depth technical discussion of networking, but rather to lay a foundation of knowledge that will assist the reader in planning and managing this technology through a better understanding of its strengths and weaknesses.

Section 1.0 first provides a general discussion of computer networking and the impetus for its development. Subsequent sections deal with basic network communications theory, the physical means of interconnecting hardware, common network operating systems, typical network applications, networking hardware, network management tools, and integration of networks in the hospital environment.

1.0 NETWORKING

In the context of this book, networking refers to the interconnection of computers and peripherals in a manner that enables these computers (and their operators) to share information, applications, and hardware resources. Early computers were standalone units – each had a separate set of functions (applications), its own data, and a specific complement of peripheral equipment (e.g., hard disk, printer, modem, tape drive, etc.). As the capability of computers increased, allowing them to handle larger and more complex tasks, access to the information they contained became critical to the work of more people. Also, as computers became a necessary tool in more jobs, the desirability of sharing expensive and otherwise underutilized peripherals increased. Large, fast hard disks, laser printers, high-speed fax/modems, CD-ROMs, and other expensive peripherals are better utilized when placed on a network to be shared by many users. Connecting computers and peripherals together into networks enables users to efficiently access the same data and to share these peripherals in a way that reduces duplication and increases utilization of resources.

Today's networked computers permit many users to simultaneously access and update information (such as accounting or patient information). Networks facilitate effective communications by enabling users to exchange a substantial amount of information with a large number of people in a short period of time [typically through the use of electronic mail or (e-mail)].

1.1 *Local Area Networks*

Local area networks (LANs) are designed to operate within limited geographical areas – cable runs are typically less than 300 meters. A LAN may serve only one department or it may connect computers on different floors, throughout a building, or over an entire corporate or hospital campus. Within the LAN's limited geographical area, data integrity can be maintained while achieving transmission rates from 10 megabits per second (Mbps) to 100 Mbps.

1.2 *Wide Area Networks*

Wide area networks (WANs) are typically comprised of a series of LANs connected over a large geographical area. These extended networks may

be linked across cities, countries, or continents. Today, the longer distances involved in WANs usually require that data be transmitted through hardware that will propagate it at slower and more reliable speeds of 1 Mbps or slower. Products currently under development provide for 100 Mbps WAN service through private or publicly established network access systems.

1.3 Internetworks

A related phenomenon has been the creation of internetworks linking multiple networks (i.e., universities, businesses, research institutions, etc.). Internetworks began to appear in the mid to late 1970s. The Internet itself evolved from an internetwork called ARPANET which linked several host computers to support and communicate research among government (primarily defense), educational, and research institutions. The Internet grew rapidly in the late 1970s with the introduction of a protocol suite called Transmission Control Protocol/Internet Protocol (TCP/IP). The Internet linked hosts and their networks of terminals to other hosts and their networks of terminals. In addition, it contained protocols for terminals, file transfer, and session management (the end-user's interaction with the operating system). The Internet marked a significant step in the spread of networks by presuming that there will be heterogeneous host hardware and operating systems.

The Internet underwent a boom in the 1980s with the introduction of PCs and PC LANs. In the 1990s, fueled by talk of the information superhighway, demand for Internet access is exploding. It has benefited from improvements in modems, communication software, and graphical user interfaces (GUIs), which simplify access and use of the Internet.

1.4 PC LANs

Local area networks for PCs began to appear as soon as the PC was debuted in the workplace (circa 1980). Like most industries in the early stages of growth, PC LANs had no standards and a multitude of competing vendors for both hardware and software. The initial purpose of PC LANs was to provide printer services and file servers to individual PCs to combat the high cost of peripheral devices—in particular, laser printers and hard disks.

Over time, PC LANs began to assume a more integrative function, supporting group activities with multiuser databases and office support software. By the late 1980s, Novell had established itself as the leader in PC LANs. Other vendors included Banyan, IBM, Microsoft, and 3COM. The latter has since exited the software market. The primary protocols for PC LANs are Ethernet, token ring, and ARCnet. ARCnet, however, is rapidly declining. For most of its history, ARCnet was kept proprietary by its inventor, the Datapoint Corporation, which stifled its evolution and improvement. Its primary advantage, cost, is no longer a factor, with Ethernet adapters now costing as little as ARCnet adapters.

1.5 Recent Trends

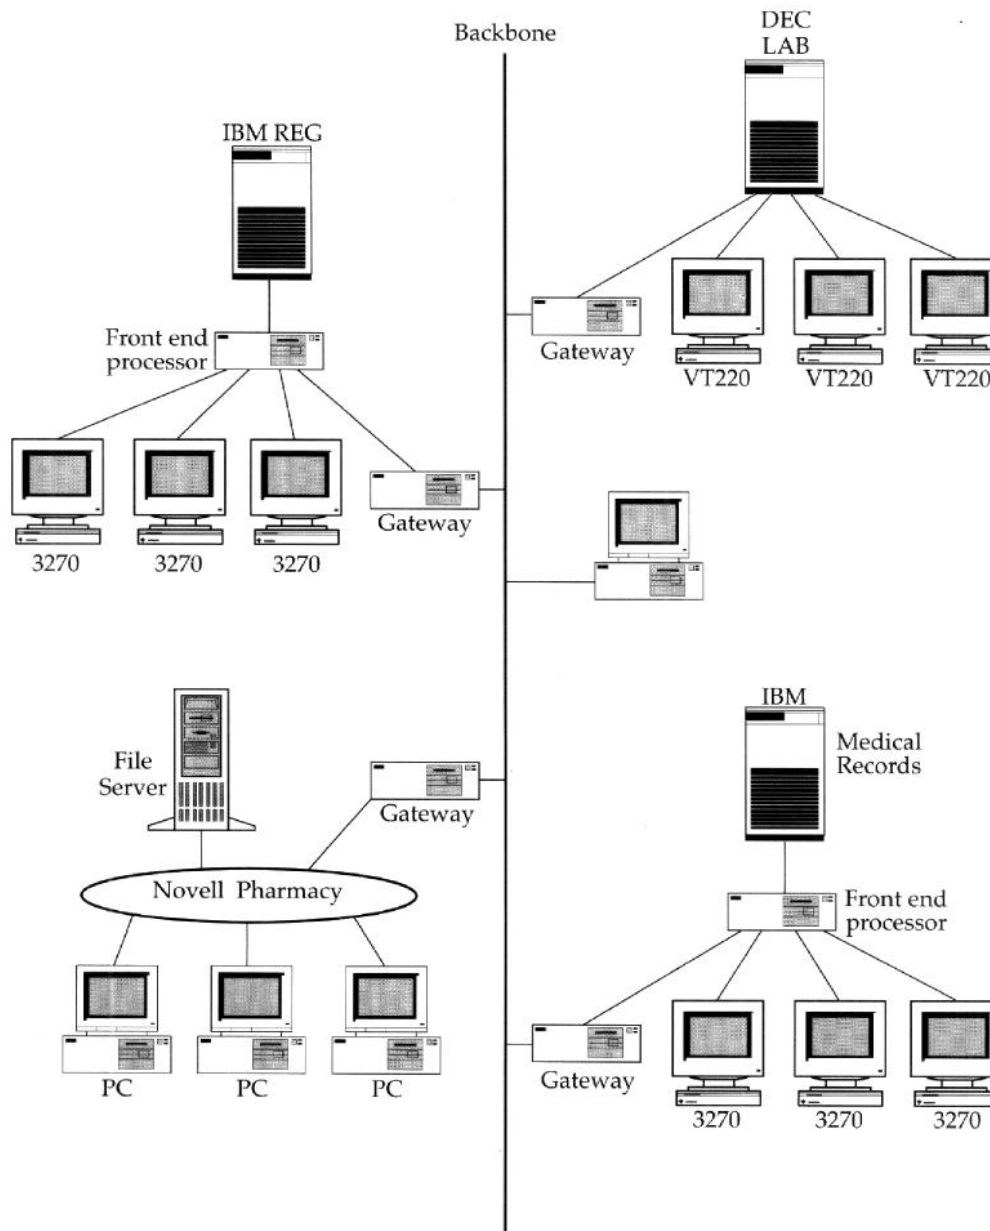
Recent trends in computer networking include: network integration, wireless networks, higher throughput, and improved tools for network management. The general result of these trends is to involve the computer in every business function, linking all functions seamlessly, without constraints imposed by location or distance.

1.5.1 Network Integration

In many organizations, the demand for integration of multiple networks continues. As each department or function has improved its intra-departmental operations through the acquisition and installation of its own LAN, organizations are coming to realize that these standalone LANs, while they provide many benefits, also have an organizational price. They reinforce departmental boundaries, impede data flows, and create discontinuities in operations and service. As a result, there is a movement to integrate LANs and large hosts on a backbone network (see Figure 1.1).

Each subnetwork (IBM, DEC, Novell) has its own conventions for cabling, plugs, network access, and addressing. In addition, the software applications for e-mail, database, and file transfer on each subnetwork are likely to be incompatible. A backbone network is a combination of hardware and software that overcome these incompatibilities. In addition, it is a common cabling system plus conversion device that sits between the backbone and the individual subnetworks. Using a variety of techniques, the backbone performs electromechanical and software translations, allowing subnetwork messages to travel throughout the entire organization.

Figure 1.1 — Integration through a backbone network.



1.5.2 Wireless LANs

The trend toward wireless LANs is related to the proliferation of cellular phones and the increasing realization that physical cables do not have to constrain the use of networks. In a LAN, moving and installing cables has always been a problem. However, wireless LANs allow workstations to be moved without the expense and delay of recabling. These LANs use a variety of technologies: microwave, radio wave, and infrared. Each has its associated costs and operating trade-offs. Wireless computing has led to the appearance of a new generation of device, the Personal Data Assistant (PDA), a combination of wireless phone and computer that can fit in a person's hand. However, the PDA has lower throughput than LANs.

1.5.3 Higher Throughput

The demand for higher transmission rates is growing for both WANs and LANs. The lack of speed (along with reliability and functionality issues) has traditionally frustrated the users and designers of WANs – an X.25 network can only transmit at 56 Kbps while a token ring LAN is capable of transmitting at 16 Mbps – almost 300 times as fast. Another issue is the reliability of the transmission. Much WAN traffic runs over older switched lines that are subject to static, further decreasing throughput. Finally, end-user functionality is severely constrained by WANs, whose operating systems do not support much functionality. The throughput factor further constrains functionality. Imaging functions on one machine cannot be shared when the lines are slow.

With WANs, increased speeds are being achieved through the use of more powerful digital transmission lines like T-1 (1.54 Mbps) and T-3 (up to 45 Mbps) and by new protocols that are able to package and send messages more efficiently (Frame Relay and Switched Multi-Megabit Data Services).

There is also a throughput issue with LANs. As more and more graphics applications are available, they are placing a severe strain on LANs. For example, a digitized X-ray can consume 16 megabytes of storage. Transmitting this image over a standard Ethernet network (10 Mbps) can take a long time and congest the network. Moreover, the addition of many more users on an Ethernet segment can tax its throughput.

LAN throughput has been increased by the introduction of new protocols and cabling schemes such as 100 Mbps Fiber Distributed Data Interface (FDDI) and 100 Mbps Ethernet (which is not yet an official standard and which is no longer Ethernet). Because of its high capacity, FDDI is being used for the backbone network in some institutions, or for LAN segments which support medical imaging workstations.

1.5.4 Network Management Tools

The expanding connection of networks creates a demand for products that permit management of all network devices in a standardized, centralized fashion. All vendors of network devices (hubs, routers, bridges, servers, etc.) include software for device management. The dominant standard is the Simple Network Management Protocol (SNMP).

1.6 Network Operating Systems

On hosts (multiuser systems) it can be difficult to clearly delimit the operating system (OS) from the network operating system (NOS). While host OSs (IBM, DEC, UNIX vendors, etc.) have NOS features built in, the original PC had no networking features in its OS (not surprising, since it was a personal computer). In a sense, Microsoft's disk operating system (DOS) was an open operating system — it allowed NOS vendors to insert their software above DOS and take it over. A shell program can sit on top of DOS and intercept network requests and execute them, or pass them on to a server.

2.0 OPEN SYSTEMS INTERCONNECTION

For purposes of clarity, it is useful to understand the role of standards and standards organizations. A standard is a written specification of the characteristics (electromechanical and functional) of the various components of a network. Without standards, all discussions of networking would degenerate into a tower of Babel, since each vendor has their own conventions and vocabulary. In order to connect two vendors' platforms, there has to be a common framework, preferably defined by a third party. This third party is the International Standards Organization (ISO), and the framework is called the Open Systems Interconnection (OSI) model.

2.1 Standards Organizations

The primary entity for the definition of worldwide standards is ISO, which is not limited to networking. ISO is located in Geneva, Switzerland, and interacts with the standards bodies of individual countries through a hierarchical network of relationships. For example, the International Electrotechnical Commission (IEC) is a part of the ISO that is directly responsible for data and voice communications. It works with the American National Standards Institute (ANSI) in the United States, which, in turn, coordinates the work of various domestic ANSI-accredited standards organizations like the International Association of Electrical and Electronics Engineers (IEEE) and the Electronic Industries Association (EIA). The International Telecommunications Union - Telecommunications Standardization Sector (ITU-TSS, formerly the CCITT) is an independent organization which often collaborates with ISO. In addition many groups have been created to produce and test products that conform to standards. They include the Corporation for Open Systems (COS) and the Open Systems Foundation (OSF).

2.1.1 International Organizations for Standardization and International Electrotechnical Committee

ISO/IEC are abbreviations for the International Organizations for Standardization and International Electrotechnical Committee. The members of the ISO/IEC are the national standards organizations, such as ANSI and the British Standards Institute (BSI). Within the ISO/IEC there are two subcommittees responsible for the OSI:

- JTC1/SC6, Telecommunications and Information Exchange between Systems
- JTC1/SC21, Information Retrieval, Transfer, and Management

JTC1/SC6 is responsible for data mobility while JTC1/SC21 is responsible for information structure.

2.1.2 International Association of Electrical and Electronic Engineers

The IEEE has developed a number of standards for communications protocols which have been adopted by the ISO. These include such well-known standards as the Ethernet Standard - IEEE 802.3 (ISO 8802/3) and the Token Ring Standard - IEEE 802.5 (ISO 8802/5).

2.1.3 International Telephone and Telegraph Consultative Committee

The International Telephone and Telegraph Consultative Committee (CCITT) has recently been renamed the ITU-TSS. The ITU-TSS is a cooperative process because the members have global coverage. For example, the ITU-TSS X.25 Level 3 protocol describes the interconnection of computers or terminals to a packet-switching network. It also produced the X.400 standard for e-mail.

2.1.4 Open System Foundation

The OSF defines the specifications for open software and also produces the software itself. OSF products include OSF1 / UNIX, an open operating system; and Distributed Computing Environment (DCE) – interoperability software using ISO standards for naming, security, and remote procedure calls. During 1993 and 1994, the organization abandoned its effort to produce software products and now is concentrating on defining the specifications.

2.1.5 Corporation for Open Systems

The COS was established by a consortium of vendors in 1986. Its purpose was to test and certify products for conformance to ISO networking standards. Organizations like OSF and COS face many obstacles. Individual vendors compete to get their convention adopted as a standard. Other vendors, whose products are widely used, refuse to submit their product for adoption as a standard. In general, the process of standards definition and testing is slow, while vendors are constantly updating existing products and creating new ones. Although many critics consider the above-mentioned obstacles to be insurmountable and helping prevent the establishment of meaningful universal standards, others

point to the undeniably huge improvement in connectivity and interoperability in the 1990s, claiming that it would not have been possible without the work of standards organizations, however suboptimal.

2.2 ***Open Systems Interconnection Model***

Before examining networks in detail, we must develop a conceptual framework. This is needed because the functions and features of a network are quite varied, and each vendor has its own vocabulary for these features. We need a Rosetta stone for translating each vendor's language back to some universal standard. This standard is the OSI model. It was developed by the ISO as a way to organize the functions of a network.

The OSI model views a network as a series of layers, as shown in Figure 2.1. Each layer is responsible for a certain piece of network functionality. When it performs this function, it passes control to another layer. The lowest layer (layer 1) provides the physical connection to the network medium. The highest layer (layer 7) provides commonality of functions at the level of the application, or end user. A summary of the layers is provided in Table 2.1.

Table 2.1 — OSI Reference Model


Network Layer	Function
Layer 7 - Application	Common or generic features of any software application: e-mail, naming, database access, file transfer.
Layer 6 - Presentation	Data formats, syntax.
Layer 5 - Session	Interface with host operating system — keeping track of the different network sessions.
Layer 4 - Transport	End-to-end integrity of the message transmitted across the network.
Layer 3 - Network	Routing of messages across and among networks.
Layer 2 - Data link	Access to the immediate network.
Layer 1 - Physical	Electromechanical characteristics of the medium and connectors.

Figure 2.1 — Layers of the OSI model.

7	PROFS
6	EBCDIC
5	MVS
4	SNA
3	SNA
2	SDLC
1	COAX

7	DECMail
6	ASCII
5	VMS
4	VMS
3	LAT
2	ASYNCH
1	RS-232

7	MHS
6	ASCII
5	NETBIOS
4	IPX
3	RIP
2	ARCNET
1	TWISTED PAIR



Each layer only interacts with the layers immediately above and below it. It provides services to the layer above, and uses the services of the layer below. This is one of the fundamental principles of structured software design — insulation of modules from each other. In theory, one can change the characteristics of a given layer (n) and only be concerned with the impact on layer n+1 and layer n-1. An example of this would be changing from RG58 coaxial cable to twisted-pair cable while maintaining the Ethernet signaling characteristics. Another example would be changing the routing protocol from Xerox Network Services (XNS) to internet packet exchange (IPX) (layer 3), keeping the transport protocol (layer 4) the same and the Ethernet protocol (layer 2) the same.

Each layer accepts data from a higher layer and wraps its protocols around it. It then passes this unit down, and it, in turn, becomes the data to be enclosed in the lower layer's protocol envelope.

To illustrate, we'll use the example of a pathologist who wants to send a report to a clinical researcher on the same campus. This is shown in Figure 2.2.

On the pathologist's side, the following things have to happen: the message has to be converted; the name of the recipient has to be common; any encryption must be done; the network must be called; the network must decide how to associate the recipient with a network and this network with a path (routing); the network OS takes care of packaging the message on one side and reassembling it on the other; the packet is further broken down into frames with the appropriate structure; the message is placed on the cable and sent to the next router; the router looks inside the packet to find the ultimate recipient, does look-ups, and passes the packet on. This is repeated for as many "hops" as are necessary. This process is reversed on the clinical researcher's side: the frames are reassembled into packets; the packets are reassembled into the message; the message is associated with a user's session and a port; decryption takes place; and the message is passed to the recipient's e-mail package. This process is greatly complicated by the presence of heterogeneous hardware and software.

2.3 *The OSI Model and Biomedical Devices*

The OSI model can provide insight into the structure of biomedical devices because the evolution of biomedical devices is similar to that of computers. Computers have evolved through the following stages: host + dumb terminals (layers 1 and 2); hosts + PCs and LANs (layers 1 to 4), where the end devices acquire microprocessors; and finally, client-server (layers 1 to 7), where there is application integration. With biomedical devices and monitors there is a similar evolution: monitors containing controllers for individual devices [RS232 and serial data link control (SDLC)]; and combining monitors into a network (layers 1 to 4). Biomedical devices range from infusion pumps to monitors (e.g., CO₂, pulse oximetry, and temperature). Originally, each type was a self-contained, standalone unit which was electromechanically controlled, rather than software controlled (see Figure 2.3).

Figure 2.2 — Routing a message from pathologist to clinical researcher.

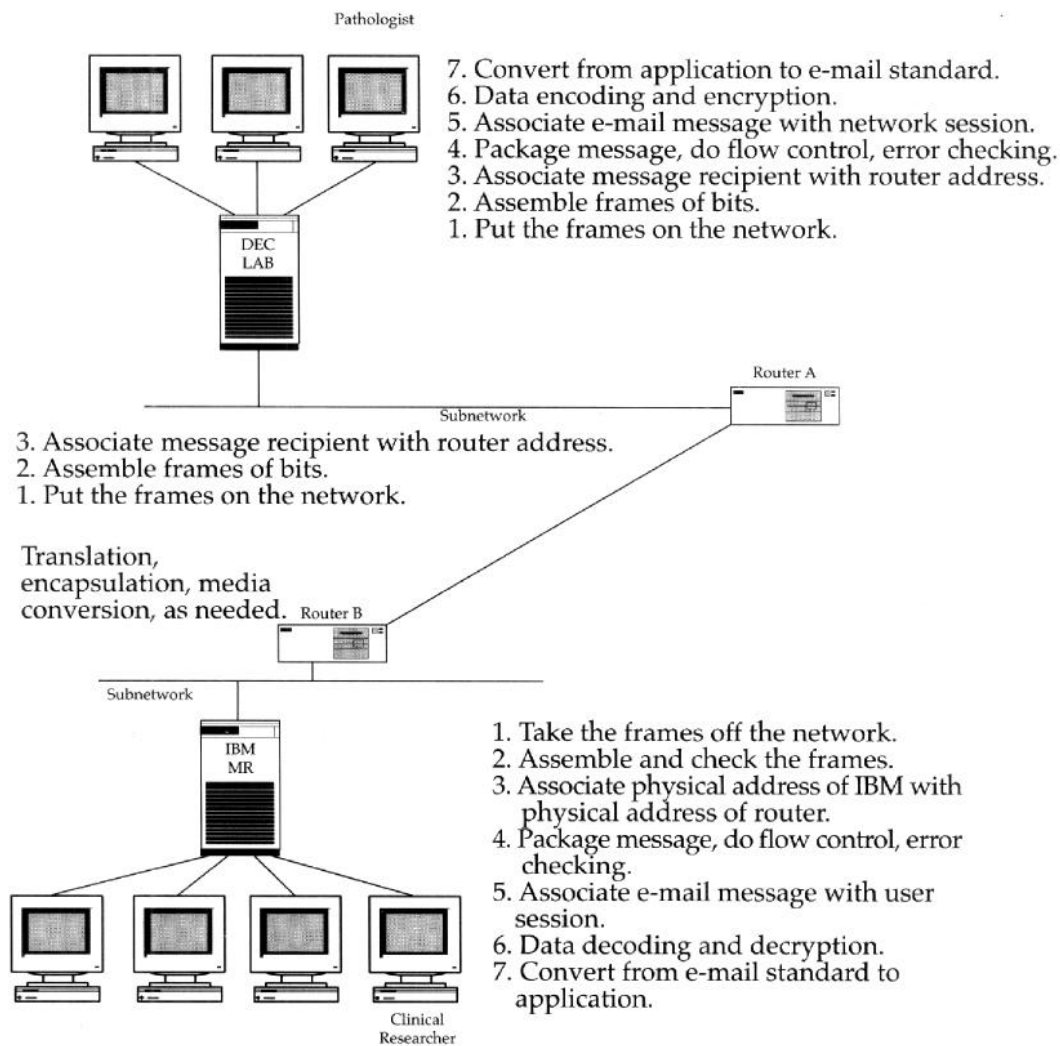



Figure 2.3 — Biomedical device to bedside monitor interface.



In Figure 2.3, each of the sensors is wired to a controller that resides in the housing of the display. This configuration is a network at OSI layers 1 and 2 only. Layer 1 specifies the electromechanical characteristics of the cable and plug; layer 2 specifies the protocol for polling the data collection devices. Over time, however, the data collection components have gotten smarter through the use of microchips, and the monitors that control them have been networked allowing remote monitoring and in some cases, remote control of devices. A human does not have to adjust them. The use of microprocessors allows digitization of the collected data; the analog data that used to be displayed on an oscilloscope is now digitized and displayed on a computer monitor; this data is available for storage. Furthermore, the digitized data is available for transfer to a clinical information system (CIS).

2.4 OSI Layer 1 - Physical

Layer 1 of the OSI model is concerned with the electromechanical characteristics of the transmission medium employed by the network. This includes voltages, number of wires, insulation, connectors, etc. Biomedical professionals are more accustomed to intrasystem communications (buses) as opposed to intersystem communications (network cables). Intrasystem communications involve multiple devices connected to the

same bus in a single machine. Intersystem communications occur when different machines communicate across a common bus or a network.

2.4.1 Topology

There is a substantial amount of confusion between network protocol (token ring, Ethernet, etc.) and network topology. The concepts of token ring and Ethernet are addressed at layer 2 of the OSI model; each is a method for gaining access to the network and transmitting data over it. It so happens that the original cable design of Ethernet was a bus, while that of token ring was a ring. While Ethernet has the logical topology of a bus, it can have multiple physical topologies and cables. Originally, Ethernet ran over thick coaxial cable with a transceiver embedded in the cable. Subsequently, RG-58 cable (thinnet) was introduced. The 10baseT, AUI, and thinnet standards were developed which allows running Ethernet over twisted-pair (telephone) wire with the wire running back to a closet. The same can be done with token ring. Figure 2.4a shows the token ring topology, while Figure 2.4b shows the same protocol, but as a star topology, in which the ring is embedded inside the MAU.


The general trend is to use wiring hubs. Thus, superficially these modern LANs look like the old systems network architecture (SNA) hierarchical networks where all communications had to be channeled upward to a front-end processor which performed network control.

The use of wiring media will be examined from four points of view: buses, point-to-point networks, LANs, and WANs. Refer to Section 4.0 for more detail.

2.5 OSI Layer 2 - Data Link

The second level of the OSI model, the data link layer, deals with accessing the immediate network (the subnetwork) and passing data frames to the adjacent node on the subnetwork. It is at this layer that it is most appropriate to talk about Ethernet or token ring. For some protocols (the IEEE 802 LAN protocols – Ethernet, token ring, FDDI, etc.), this layer is actually divided into two distinct sublayers, the Logical Link Control (LLC) and the Media Access Control (MAC). The LLC is independent of the medium employed, while the MAC has a separate format for each of the media. It is possible for multiple, higher-level protocols to coexist on the same layer 2. For example, some workstations could be running No-

Figure 2.4 — Token ring: (a) ring topology; (b) star topology.



vell IPX, while others are running IP; their differences are transparent to layer 2. IPX packets would only be recognized and received by workstations running Novell; IP packets would be ignored. Similarly, TCP/IP workstations would ignore IPX packets.

2.6 ***OSI Layer 3 - Network***

The third level of the OSI model, the network layer, ties different subnetworks together by handling the routing of messages across networks. Network layer software uses routing tables and a variety of routing algorithms. It is at this level that many problems occur because many vendor-specific protocols (i.e., NetBIOS, SNA, and local area transport) were developed without provision for internetworks, and thus, their layer 3 protocols are not routable.

There are different routing approaches: spanning tree algorithm, source routing, transparent routing, source transparent routing, static routing, and dynamic routing. With the trend to create enterprise networks and internetworks, a common technique is to perform tunneling, whereby one transport protocol is wrapped within another; see Section 2.6.4.

The international standards for layers 1 and 2 are well established and widely accepted. New standards are continually being proposed as technology evolves – a case in point is the emergence of asynchronous transfer method (ATM). Unfortunately, the standards for layers 1 and 2 say nothing about how messages should be routed. There are well over a dozen routing algorithms. For purposes of discussion it is useful to divide network protocols into 1) international standards, 2) de facto standards, and 3) vendor conventions.

2.6.1 International Standards

OSI has two standards: Intermediate System-to-Intermediate System (IS-IS) and End System-to-Intermediate System (ES-IS). The former is used by routers to evaluate the best path to a given node on the network; the latter is used by an end-user node to inform neighbors of its presence. The OSI protocols are not widely implemented for two reasons:

- They require a lot of system overhead for set-up.
- While standards work was being focused on layers 1 and 2, vendors were actively working on conventions for layers 3 and above, outpacing the standards groups.

The X.25 protocol was developed in the late 1970s by the CCITT as a method for connecting terminals and computers over a wide area network – usually a public data network. The X.25 is a packet-switched network. The X.25 connections, called virtual circuits, may be established temporarily, such as a dial-up connection, or permanently, like a leased line. The X.25 network supports multiple-user sessions over the same physical link. A number of virtual circuits are supported on a single physical connection by the individual identification of packets across the network. As a protocol, X.25 spans layers 2 and 3. X.25 has been used to connect local SNA sites across a wide geographical area. It has also been used to connect local TCP/IP sites. The X.25 is being phased out by emerging standards such as frame relay.

2.6.2 De Facto Standards

Given the slow adoption rate of the OSI standard, the routing information protocol (RIP) convention in TCP/IP has become the de facto standard for internetworking. Routing information protocol is a routing algorithm based on counting hops (intermediate nodes between sender and addressee). The sending node decides whether the destination node is on the same subnetwork. If not, the sending node forwards the message to a router, which, in turn, determines the best route. The number of hops is limited to 15. Open Shortest Path First (OSPF) has been proposed as a solution to this limitation.

2.6.3 Vendor Conventions

Given the slow promulgation of OSI standards, many vendors have promoted their own conventions. This section examines the issue from three different computing levels: mainframe (IBM), minicomputer (DEC), and PCs (Novell).

2.6.3.1 IBM Systems Network Architecture

Systems network architecture is IBM's traditional networking architecture. Conceived in the 1970s, SNA was originally designed to support terminals and printers connected to mainframes through controllers and front-end processors. SNA assumed a hierarchical structure: the end-user devices were dumb terminals, not central processing units (CPUs) in their own right; and all communications were sent upward. Emphasis is on connection-oriented sessions. The end-user establishes a mutually agreed upon connection with the host and there is synchronous communication for the duration of that session. The terminal regularly sends a confirmation message that the user is still logged on. If the host does not receive this message within a predetermined amount of time, the host cancels the session. The process of setting up and breaking down a session was complex and lengthy, but transactions were processed quickly. The emphasis was less on flexibility and more on guaranteed response time (e.g., bank terminal networks and reservation systems). Over time SNA has evolved, gradually assuming more and more features of LANs. First, the concept of peer-to-peer IBM devices was introduced; this was followed by LAN concepts; finally, IBM has announced APPN,

a convention that allows heterogeneous peer-to-peer networking over SNA networks.

2.6.3.2 DEC Local Area Transport

Local area transport (LAT) is a convention developed by DEC to accommodate asynchronous terminals attached to terminal servers on an Ethernet. LAT offloads terminal handling from the minicomputer and places these functions in the terminal server, which then communicates with the host through a LAN protocol such as TCP/IP. Software in the terminal server also allows routing of LAT, which is otherwise not routable.

2.6.3.3 Novell Internet Packet Exchange/Sequenced Packet Exchange

IPX is a Novell convention, a modification of XNS. It is frequently mentioned in conjunction with Sequenced Packet Exchange (SPX), which is a layer 4 and above protocol, enabling program-to-program communication, error correction, checkpointing, and flow control. Initially, Novell was used for small workgroup LANs. Over time, Novell has grown and is achieving dominance in the LAN marketplace and being used for company-wide LANs, including WANs. Unfortunately, the IPX protocol is not suited for use in WANs. Novell recognizes this and is moving away from IPX and toward TCP/IP.

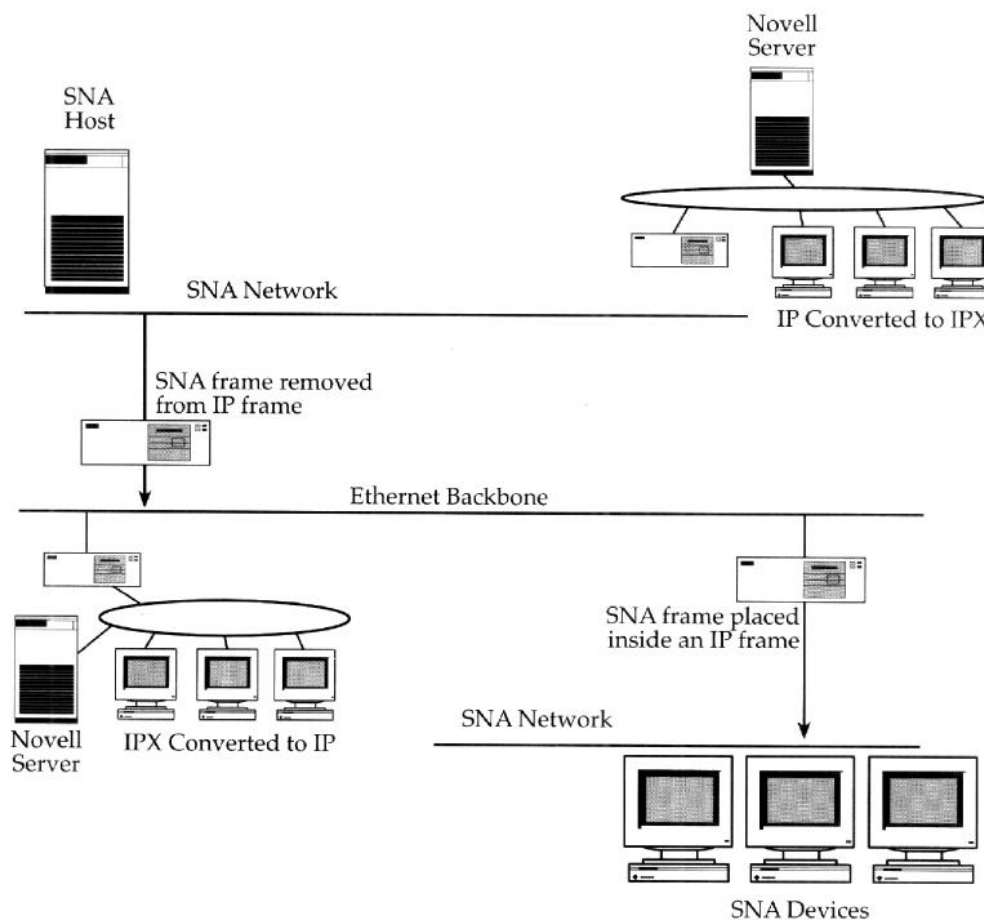
2.6.4 Tunneling

Assuming that a company is unable/unwilling to replace conflicting networks with a single standard, there are two ways to overcome routing incompatibilities: translation and tunneling. In translation, the frames of one network are converted into the frames of another network, through the use of a multiprotocol router. Figure 2.5 illustrates two Novell workgroups located at opposite ends of a campus and are interconnected through the TCP/IP backbone. On the sending side IPX frames are converted to IP. On the receiving side, these IP frames are converted back to IPX. However, with tunneling there is no translation. Rather, the frame of one network type is encased within a layer 3 frame of another network. In Figure 2.5 several SNA terminals are located at one end of the campus and they need to be connected to another SNA host on the other side of the campus. The already existing backbone is

selected as the medium. As they pass through the backbone router, the SNA frames are enclosed within IP frames, with no conversion of their structure. At the other end, the IP envelope is stripped away, and the SNA frame is sent to the host.

Tunneling and conversion have advantages and disadvantages. Both require additional overhead in the processing of frames. The delay associated with tunneling is particularly critical for SNA transactions, which are time sensitive. On the other hand, tunneling is a less complex operation than translation, which can introduce errors when functionality is lost in the translation process.

Figure 2.5 — Tunneling and conversion.



2.7 OSI Layer 4 - Transport

The transport layer handles the end-to-end integrity of the message. This includes flow control (timeouts), checkpoints, sequencing of packets, checking for lost packets, error checking, etc. The type and amount of error checking performed at this layer depends on the reliability of the lower layers (1 and 2). The greater the reliability of the medium and the more error checking done at layer 2, the less need for error checking at higher levels. Reliability factors include network performance, OS and application tuning, buffer size, the number of hops and delays, the bandwidth available, and the quality of the network (the amount of noise).

As with the network layer, it is useful to structure the discussion in terms of international standards and de facto standards. However, vendor implementations of the transport layer show much less conformance to international standards. Many of the functions specified in the OSI model are commingled with the host OS functions and application features, blurring the clean “layering” specified by OSI. This confusion is one of the contributing factors in the growth of software called middleware (see Section 2.8.1).

2.7.1 International Standards

OSI has defined five different classes of service for the transport layer (TP0 through TP4). While they all provide the basic services of connection establishment, message segmentation, and message reassembly, TP2 adds multiplexing, and TP4 adds error detection and recovery.

2.7.2 De Facto Standards - TCP/IP

TCP/IP provides roughly the same types of services as OSI. It is more widely used. It supports both connection-oriented messages (TCP) as well as connectionless messages – user datagram protocol (UDP). The TCP protocol specifies error checking and correction, while UDP simply performs “best effort” delivery.

2.7.3 Vendor Conventions - SPX

Novell uses SPX for its layer 4 protocol. However, SPX is more suited to a LAN environment than a WAN environment. As Novell extends its reach to include the whole enterprise, meaning geographical dispersion, it will further unbundle SPX from its Netware core and replace it with TCP.

2.8 OSI Layer 5 - Session

The fifth level of the OSI model, the session layer, sets up and terminates application-to-application dialogs, associates network messages with end-user sessions, and sessions with ports. It represents the handoff between the NOS and the node's own OS. Frequently, it is difficult to distinguish between OS and the NOS features. UNIX, for example is on the operating system which has networking features built in.

It is at this point that it becomes difficult to separate the OS from the NOS. The OSI standards for the session layer have not been widely adopted. Instead, vendor conventions dominate, with OSI layer 5 functions commingled with operating systems and applications. For this reason, the discussion of layer 5 and layer 6 standards are not divided into international and de facto standards.

2.8.1 Middleware

As has been mentioned previously, layers 1 to 3 have been clearly defined and widely promulgated. However, there remains a gap between the application and layer 3 – caused by the slow spread of OSI standards and the multiple competing vendor conventions. This causes severe problems for the application developer, who has to master the idiosyncrasies of operating systems and networks as they concern the functions specified in OSI layers 4 to 7. The application developer must know much more than the structure of their specific application. They must be able to construct calls to OS and NOS programs, send the correct parameters, and understand the data returned. These functions are quite diverse and include end-to-end synchronization, interface to operating system, session coordination, network calls, compression, encryption, syntax, database calls, e-mail calls, etc.

The answer to this is middleware – a new breed of software that insulates the application developer from the complexity of the OS and NOS. Using a simplified application programming interface (API), the developer makes a call to the middleware, which handles all the other calls to network modules, program libraries, etc. A further benefit of middleware is that it creates greater portability of applications, allowing them to be moved to a variety of OSs and NOSs.

2.9 ***OSI Layer 6 - Presentation***

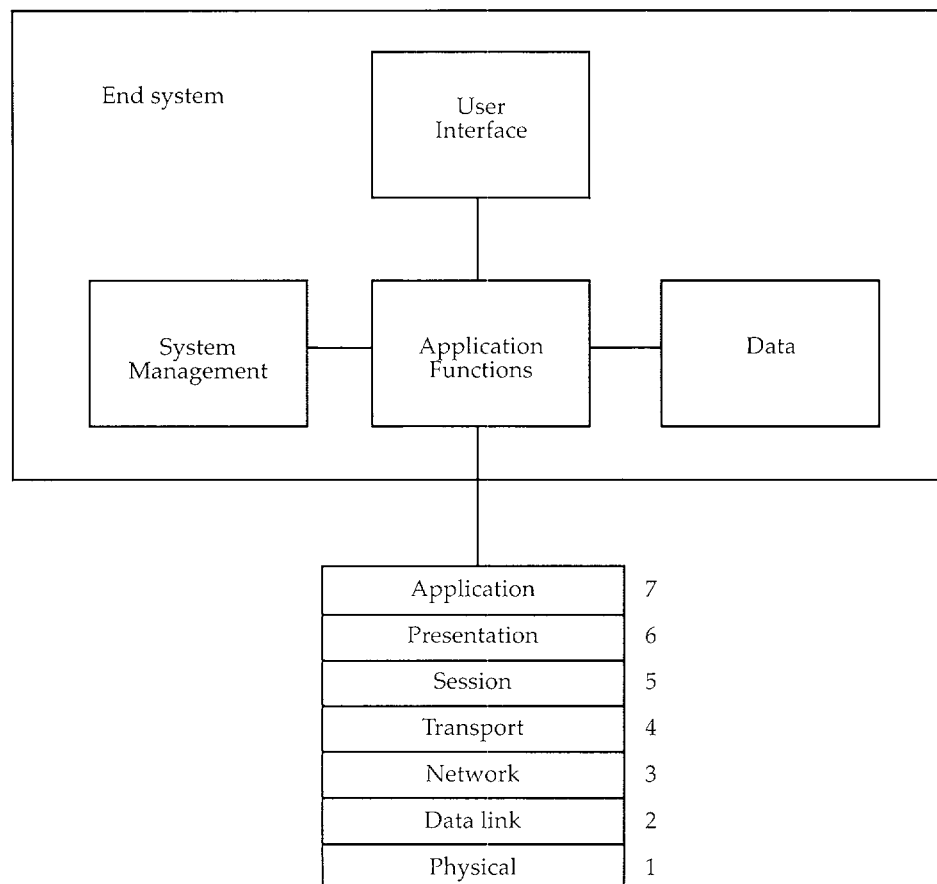
This layer deals with the structure of the data that is being transmitted – its record structure, encryption, and compression. With regard to data structure, layer 6 specifies how applications will indicate the structure of the data that is being transmitted – records, record types, fields, sub-fields, data types, etc. For example, an HL7 admission message has to specify how its records (header, admission, trailer) are delimited, what their constituent fields are, the data types of the fields, and how they are delimited. See Section 8.1 for further discussion of HL7.

The OSI standard for data syntax is abstract syntax notation (ASN.1). Once again, there is little evidence that the OSI services are being consistently implemented by vendors in a separate software layer. Rather, they are subsumed by applications and operating systems in a variety of ways. Such inconsistency further explains why application developers are so eager to turn this over to middleware.

2.10 ***OSI Layer 7 - Application***

The highest level of the OSI model, the application layer, represents the network's interface to end-user applications (e.g., patient registration, medical records, e-mail, nursing notes, etc.). As OSI was originally envisioned, the operating systems and applications were supposed to be treated as black boxes whose functionality was totally invisible to the network. While oblivious to the functionality above it, layer 7 also contained no functionality (services) – it did nothing beyond accepting the data elements that needed to be passed by the functions of a given entity in the black box. For example, the standard for e-mail specified which fields needed to be passed, what their meaning was, and their relative position (e.g., the fields for message sender, message recipient, message priority, message cc:, receipts, etc.). However, there is a fallacy in this logic. It is impossible to neatly separate the functionality from the interface. However, the very definition of the types of fields that must be passed by the black box to use the e-mail function says something about the structure of the application which is supposed to be in the black box. And so, over time, the concept of layer 7 seems to be evolving to include greater amounts of functionality and a wider range of functionality not foreseen in the original design of layer 7. This includes operating system portability, user interface (e.g., Microsoft Windows), system administration and database access (Figure 2.6).

Figure 2.6 — Application features of layer 7.



This black box approach was greatly conditioned by the technological environment in which the original OSI designers worked. Their world was dominated by hosts and dumb terminals. WANs dominated; there were few LANs. There were no PCs and few end user tools. User interfaces and sophisticated system management were not an issue. End-user functionality was limited to the exchange of e-mail and simple documents, file transfer, remote procedure calls, and terminal emulation. The appearance of PCs and LANs changed that. The amount of new-found functionality in the black box created the need to insert greater functionality into layer 7, which now provides services for database access, naming, complex document exchange, e-mail, complex docu-

ments, security, system administration, user interface, electronic data interchange, and new types of remote procedure calls using object-oriented technology.

While reading the next section of the book, the reader should keep in mind that the purpose of layer 7 is to ensure that end users and their applications can interact in a way that mirrors the structure of the business as closely as possible.

2.10.1 Naming

Determining the names of the entities with which one wants to communicate across a network of networks is not a trivial issue. Different networks may have different local names for a person or device, and it can be very difficult to know in advance what something is called, particularly if it is geographically remote. Therefore, it is important to have a way of disambiguating every object. This is addressed in the X.500 standard, which allows an application or a user to look up the name of an object by querying a database using an alias or the value of an attribute, such as street address or telephone number. Every object is uniquely identified within a tree-structured directory whose levels include country, organization, organizational unit, person, and role. Objects like devices have a different path on the tree, specifying applications and devices. This tree-like structure is similar to the Internet convention "john@gigatech.com."

This type of directory should not be confused with routing tables, which relate a user's logical name to his or her physical address on the network. The question of global naming has become particularly acute in the last several years as vendors and users attempt to integrate multiple e-mail packages across an enterprise. The literature is replete with examples of projects that have failed due to the inability of individual vendors to accommodate global naming.

The X.500 has not been widely adopted until now. However, the recent growth in the acceptance of OSFs distributed computing environment, which has X.500 as one of its components, should stimulate its adoption.

2.10.2 Electronic Mail

The most common convention for e-mail on the Internet is simple mail transfer protocol (SMTP) from the TCP/IP suite of upper-level protocols. SMTP is a store-and-forward model, that is, when a mail message is

transmitted by the sender, it is relayed to an intermediate computer where it is stored until it can be forwarded to the recipient's computer. When the message arrives at the recipient's computer, it is placed in a queue and later moved to the recipient's mailbox storage area. This method is preferred to transmission through a direct connection; it requires less bandwidth; it accommodates the security, timing, and capacity constraints of the recipient's system; and it permits the use of gateways to other e-mail systems. One of the limitations of SMTP is that, until recently, it only supported the transfer of single-message ASCII text. However, recent extensions support the transmission of multipart messages that can contain documents with word-processing formats, images, and voice.

Over the next few years, SMTP will probably give way to X.400, the OSI standard for e-mail. Like SMTP, X.400 defines a general store-and-forward message service. However, X.400 has a number of additional features such as support for international alphabets and support for a wide variety of data types (binary, images, digitized voice). A further advantage of X.400 is that its design allows it to be used as an e-mail gateway to connect the many proprietary mail systems in use.

Efforts to connect multiple vendors' e-mail products into an integrated enterprise-wide e-mail system have proven less than successful. Each vendor has its own conventions for what fields can be passed (e.g., cc:, priority, attachments, address, etc.), for message syntax, and other features such as encryption. The leading vendors of e-mail products (Novell, Lotus, and Microsoft) have each proposed their own convention for adoption as the standard. To translate from one vendor convention to another, companies construct e-mail gateways. However, in the process of translation several problems occur: the conversion is done correctly and a portion of the message is omitted or distorted, or conversion is impossible because vendor A does not support all the features of vendor B. These gateways are further complicated by the fact that they frequently require a separate interface between each pair of e-mail packages. It would be simpler to use X.400 as the gateway.

2.10.3 Complex Document Exchange

Traditional e-mail assumes the exchange of unstructured text plus several fields. However, it is possible to construct a document that consists of tables, graphic data, traditional data, and even voice and video data.

2.10.4 Remote Procedure Call

A remote procedure call makes it possible to execute a program that is located on another remote computer. This approach is used when the remote computer has special computational capacity needed by the local computer or has data that needs to be shared by many different users.

The second case can take many different forms: file server, database server, name server (for e-mail naming), message server (for store-and-forward), etc. A remote procedure capability is a *sine qua non* for network connectivity, and is used by almost all other layer 7 services. A remote procedure call looks like a call to a subroutine on the local computer; however, the subroutine is located on the remote computer.

This remote procedure call is named RPC in the TCP/IP suite. It can run in connection, oriented mode or connectionless mode. Connectionless mode has less network overhead. It can also run synchronously (a response is required from the remote system before the local system can continue) or asynchronously (the local system may continue processing before it receives a response from the remote system). The OSI protocol for remote procedures is called remote operations service element (ROSE).

2.10.5 File Exchange

In a networked environment, a common activity is the copying of files from one computer system to another. Superficially, this may seem like a simple operation, but each network vendor has their own conventions for the naming and accessing of files, for the traversing of directories, and for the types of data that can be represented in files. The TCP/IP world uses file transfer protocol (FTP), which is rather basic, but adequate for many situations involving simple file transfer – the file itself (block versus stream transmission mode, error recovery and restart, etc.), viewing directories, and utilities for renaming and deleting files.

The OSI standard is called file-transfer access and management (FTAM). It has much greater functionality than FTP, including: the ability to navigate a file, more complex file structures, the ability to transmit individual records, etc. FTAM is not widely implemented because of the large overhead it requires.

Another popular protocol is network file system (NFS). NFS was originally developed by SUN Microsystems, but SUN has published it as a TCP/IP protocol. NFS makes remote files and directories appear to be part of the local system.

2.10.6 Remote Database Access

A database, regardless of whether it is relational or not, provides a higher level of functionality and complexity than a file management system. A database uses the lower level file management system of the host and provides a higher level of logical access to the data, providing a view of the data that combines multiple files, but at the same time shielding the programmer and the end user from the physical details of the underlying files. Each hospital really has a single logical model which reflects the structural relationships among patients, physicians, procedures, diagnoses, etc.; however, the data is usually distributed throughout numerous physical systems, which may use different operating systems and database management systems. Nevertheless, there is a need to manage the data in a logically integrated fashion. The recent advances in networks and relational databases have, for the first time, made it possible to align the physical systems with the logic of the hospital. OSI has two standards in this regard – remote database access (RDA) for the opening and closing of the remote database, and structured data query language (SQL) for the actual manipulation of the data. RDA makes use of ISO's ROSE.

Remote database access has much overhead, and its adoption has been slow. Many of the database management system (DBMS) vendors have developed their proprietary middleware solutions to provide RDA. In the past several years, Microsoft has started to lobby for adoption of its open data base connection (ODBC) as a standard for RDA.

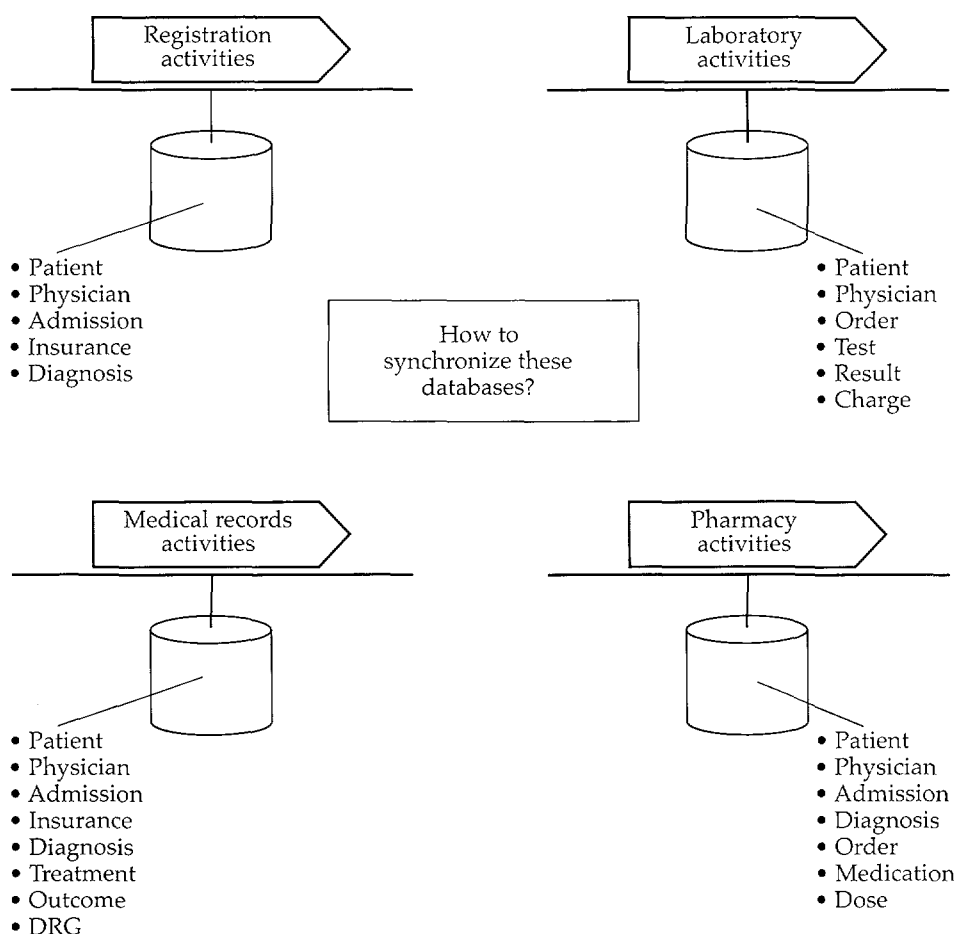
2.10.7 Electronic Data Interchange

Most people, when they think of electronic data interchange (EDI), conjure up an image of dialing into a supplier and sending an electronic purchase order, or sending electronic debits and credits to a bank. While these are valid examples of EDI with external organizations, most EDI is done internally among the information systems of various hospital departments. The area of data interchange has grown in the past several years because improvements in networking and databases have made it possible to synchronize multiple databases across the organization or within a department.

Electronic data interchange is a rational approach to resolving problems of system interfaces. Information systems contain a great deal of duplicate data. In the diverse information systems of a hospital, there are

many classes of data that are common to them all (e.g., patient, doctor, charge). However, each system has a set of core data that distinguishes it (lab has test and result; pharmacy has dose and allergy), but these data classes are useless unless they are related to the general classes of patient, doctor, order, etc. Traditional EDI involved a tape produced from one system and sent to another where it was loaded in batch after some delay. The new approach is to construct a model of how the data fits together, and determine those key events that trigger a synchronizing transaction to be sent to all systems that have an interest in that data (see Figure 2.7).

Figure 2.7 — Multiple databases needing synchronization.



Standards for data interchange exist in every industry. Manufacturing has manufacturing automation protocol (MAP) and technical office protocol (TOP). Banking and retail also have their standards. Healthcare has HL7, Medical Data Interchange Committee (MEDIX), American Standard for Testing and Materials (ASTM), ACR-NEMA, and MIB.

Electronic data interchange permits just-in-time inventory control, electronic payments, shared product specs, and a variety of interorganizational linkages. The significance of EDI is not usually fully appreciated — it is not just the mechanical exchange of data, it is also a technique for integrating the activities of the organization and can have a profound effect on quality of service.

3.0 NETWORK INTERFACES

The network interface is the physical connection between the computer and the network cabling system. There are a variety of network interfaces available, and the specific type determines the method used to send and receive data, rate of transmission, size and makeup of data packets (or frames), cable access method, network topologies, and the supported cable types.

Typically, the network interface is an adapter or network interface card (NIC) that attaches to the computer bus via an available expansion slot in the computer. Nearly all microcomputers have expansion slots to accommodate adapters that will provide the computer with additional features. On an IBM-compatible computer system these slots (or connections to the computer's bus) are likely to meet one of the following industry standards (listed here in an order relative to the age of the standard):

- 8-bit for Intel 8086/8088 XT type bus.
- 16-bit for Intel 80286 AT type bus and later (also called Industry Standard Architecture or ISA).
- 8-bit or 16-bit peripheral computer memory card (PCMCIA) most commonly found in notebook computers.
- 32-bit for IBM's microchannel bus (MCA).
- 32-bit for Enhanced Industry Standard Architecture (EISA) bus.
- 32-bit Video Electronics Standards Association (VESA) bus.
- 32- or 64-bit for Intel's Peripheral Component Interconnect (PCI) bus.

The NIC installed in the computer should be designed for the widest bus connection the computer is capable of supporting. NICs designed for the 64-bit bus or 32-bit bus will generally outperform those designed for a 16-bit bus. While representing a small percentage of the total installed local area network (LAN) base, there are also "zero-slot" LANs which do not require the installation of network hardware in one of the expansion slots on each computer, but instead make use of the computer's existing serial (RS232) or parallel port. With either a NIC or zero-slot-based LAN, some type of cable is usually run between computers (or between each computer and a hub / concentrator). A network operating system must also be loaded and configured on a computer that acts as a server or host.

3.1 *Cabling*

Various types of cable are used to interconnect PCs. The cable's outer jacket is usually polyvinyl chloride (PVC). In installations requiring fire-resistant material, Teflon-coated or plenum cable is used.

3.1.1 Coaxial Cable

Coaxial cable (coax) is constructed with a core copper wire covered by plastic insulation which is surrounded by a metallic foil or woven mesh copper shield, followed by yet another layer of insulation as an outer jacket. The outer conductor both shields the inner conductor against external interference and reduces the effects of radiated electrical noise generated by signals passing through the inner conductor. The cable's relative lack of susceptibility to radio-frequency interference (RFI) or electromagnetic interference (EMI) makes it well suited as an inexpensive media for use in applications involving high data transmission rates and long distance runs.

3.1.2 Twisted-Pair Cable

Twisted-pair cable consists of one or more pairs of wire twisted together (at a fixed number of twists per foot) within an insulated sheath. The twisting provides a mutual canceling effect against radiated interference. Twisted-pair cable includes two major categories described below.

3.1.2.1 Unshielded Twisted-Pair

As the name implies, unshielded twisted-pair (UTP) cable consists only of one or more pairs of cable twisted together within an unshielded insulated jacket.

3.1.2.2 Shielded Twisted-Pair

Shielded twisted-pair (STP) cable includes a wire mesh or foil surrounding the twisted pairs. This shielding further reduces both the amounts of interference generated by signals passing through the cable and the cable's susceptibility to externally generated interference. As a result, the STP cable can more easily transmit signals at higher rates and over longer distances than UTP.

3.1.3 Fiber-Optic Cable

Fiber-optic cable consists of thin strands of glass covered by a protective sheath of material (such as Kevlar). Since signals are transmitted through fiber-optic cable as pulses of light generated by small lasers or light emitting diodes (LEDs), they generate no interference and are immune to interference generated by other sources. Fiber-optic cable is, therefore, ideally suited for transmitting large amounts of data at high speeds over long distances.

Fiber Distributed Data Interface (FDDI) is a standard developed by the American National Standards Institute (ANSI). It provides for the transmission of data at 100 megabits per second (Mbps) simultaneously along two physical rings in both directions.

3.1.4 Wireless Connections

Some network environments can significantly benefit when computer workstations and peripherals are not constrained by a cable tether. While data transmitted via wireless technologies usually cannot be moved in the same volume or at the same high speeds of a good cable system, wireless connections do permit relatively easy movement or relocation of network components.

3.1.4.1 Infrared Transceivers

Infrared transceivers provide one means of wireless data propagation over some LAN segments. Because infrared light has the ability to bounce off walls, ceilings, and other structures, transceivers do not have to be located within line-of-sight of each other. However, since infrared light cannot pass through walls or floors, infrared transceivers are limited to use within a room.

3.1.4.2 Radio Signals

Radio signals can pass between floors and windows, making wireless radio connections between LAN segments possible over longer distances. However, most radio signals are subject to interference (as well as generating interference themselves) and may be susceptible to degradation in strength as distances increase. Transmitting digital radio signals and the use of a distributed array of antennas (as in the digital cellular systems) can help minimize the problems of interference and signal degradation.

3.1.4.3 Microwave Signals

Microwave signals can pass over long distances but transceivers must remain in line of sight of each other. Microwave is particularly well suited to transmissions between buildings where cabling is not feasible (e.g., between high-rise office buildings) or across greater geographical distances via satellites.

3.2 Topologies

The physical layout of a LAN cabling system (i.e., the manner in which cable interconnects network components) is its topology. The design of the network topology affects cost, ease of maintenance, reliability, and imperviousness to interruption. Network topologies are generally constructed of one or more of the basic elements described in the following sections.

3.2.1 Bus Topology

A bus topology (also called daisy-chain or linear) connects network components in a line, one after the other (Figure 3.1). Each network component taps into the cable as it snakes its way along from one component to another. This design requires that signals generated by one

network component and destined for another must be broadcast to all devices on the bus. A bus is generally the simplest and least expensive topology to use in small networks. Because a bus network relies on a common data highway, a malfunctioning node simply ceases to communicate; it doesn't disrupt operation as it might on a ring network in which messages are passed from one node to the next.

3.2.2 Star

The star topology (also called a hub) connects each network component to a wiring center to propagate signals between components (Figure 3.2). This cabling topology looks like a star, where "rays" extend out as cable segments to each component.

The star topology offers several advantages over the bus topology for larger networks. It is generally easier to add or relocate network components in the star wiring scheme since only one cable segment must be added or moved. Also, identification and isolation of problem cable segments or components is easy. A suspect segment or component can be disconnected without disrupting the entire network. A star topology also has the potential for accommodating hardware that can efficiently manage data traffic along the cable. While network traffic is typically generated by one component and seen by all components including the intended recipient, the segmented nature of a star topology permits the use of hardware (such as switches, routers, and gateways) that can selectively pass information to the branch where the intended recipient resides.

3.2.3 Ring

A ring topology connects each network component to both the components before and after it on the ring (Figure 3.3). Signals are typically passed along from one component to the next in one direction along the ring rather than broadcast (as is usually the case in bus and star topology networks).

Figure 3.1 — Network components: bus configuration.




Figure 3.2 — Network components: star configuration.





Figure 3.3 — Network components: ring configuration.



3.3 **Wiring Centers**

In star and ring topologies, network components are connected via cable segments back to a wiring center. Wiring centers may be relatively simple devices for receiving and relaying signals between components or they may be more complex devices that interconnect a variety of cabling systems and provide the ability to monitor and manage network performance. There are four basic terms used to describe wiring centers: hubs, concentrators, multistation access units (MAUs), and repeaters.

3.3.1 Hubs

A hub is typically a simple device housed in a cabinet which connects nodes via a specific type of cabling. Hubs are best suited for small networks (i.e., five to 25 nodes).

3.3.2 Concentrators

A concentrator is a more complex device that may have multiple modules which provide the ability to connect nodes via different cabling, to act as routers or bridges, and to provide network management services.

3.3.3 Multistation Access Units

The term MAU is typically used to refer to a token ring wiring center. It functions similar to hubs and concentrators.

3.3.4 Repeaters

A repeater is a device often used in conjunction with wiring centers to extend network cabling segments over longer distances. The repeater receives the network signal from one cable segment, amplifies it, and retransmits it over another segment.

3.4 LAN Standards/Architectures

Network interfaces are designed to conform with a specific LAN standard. This ensures they will be able to communicate with other network elements and interfaces from other manufacturers. The LAN standard to which the network interface subscribes specifies packet (or frame) structure, cable access method, signal strengths, cable types, and allowable cabling distances. These standards have either been developed and published by one manufacturer (or several working together) or by the Institute of Electronic and Electrical Engineers (IEEE).

3.4.1 IEEE 802.3 - Ethernet

The Ethernet specification was jointly published by Digital Equipment Corp., Intel, and Xerox in 1980, with a revision published in 1982. In 1985, the 802 subcommittee of IEEE approved a modified version of the Ethernet specification as a standard. This specification, which carries the

designation of 802.3, is commonly but erroneously referred to as Ethernet. While the 802.3 specification incorporates most elements of the earlier Ethernet specification, it also incorporates changes in the basic structure of the data packets being transmitted, thus making the two standards incompatible.

Both Ethernet and 802.3 make use of the same cabling, have throughput specifications of 10 Mbps, and utilize a cable access method called carrier sense multiple access with collision detection (CSMA/CD) where a network component listens to the cable before broadcasting a signal. If no other network component is broadcasting on the cable, it transmits its data; otherwise it waits and then checks again to see if the cable is quiet. When two network components simultaneously transmit data, a collision occurs and the higher signal level it generates is detected by the transmitting components. Each station backs off for a randomly determined period before attempting to rebroadcast.

3.4.1.1 10Base5 Standard

10Base5 is an IEEE 802.3 standard for the use of thick 50-ohm coax (also called thicknet). 10Base5 cable is sometimes called either yellow cable (its usual color unless the orange plenum cable is used) or garden hose cable (because of its stiffness). In a bus topology, network components are connected along a length of 10Base5 cable that may be as long as 500 meters (with a 50-ohm terminating resistor on either end). Each component connects to this cable via an attached transceiver and a shielded attached unit interface (AUI) cable that may be up to 50 meters in length (Figure 3.4).

This cabling permits propagation of signals at 10 Mbps over long distances. However, the cable is relatively thick, hard to install, and expensive. It also shares the disadvantage of other bus topologies in that a break in a 10Base5 cable segment can completely disrupt network operations.

3.4.1.2 10Base2 Standard

10Base2 is an IEEE 802.3 standard utilizing 50-ohm RG-58/A-AU coax cable (also called thinnet) to connect network components over a distance of up to 185 meters in a bus topology. The network interface of each component is attached to the cable via a T-connector (Figure 3.5).

10Base2 cable is lighter, easier to handle, and less expensive than 10Base5. Like 10Base5, it also transmits signals at 10 Mbps. However, its

range is more limited than 10Base5 and its bus topology also leaves the network susceptible to interruption should a break occur in the cable.

Figure 3.4 — Ethernet 10Base5 thicknet cabling configuration.




Figure 3.5 — Ethernet 10Base2 thinnet cabling configuration.





Figure 3.6 — Network components: Ethernet 10BaseT unshielded twisted-pair cabling.



3.4.1.3 10BaseT Standard

10BaseT is the IEEE 802.3 standard developed for the use of UTP cable. A 10BaseT network is configured in a star topology where cable segments radiate out from a central hub or concentrator to each network component (Figure 3.6).

The UTP cable used in 10BaseT networks is commonly employed in telephone wiring and is relatively inexpensive and easy to work with. As with 10Base2 and 10Base5, 10BaseT cable passes signals at 10 Mbps. While individual cable segments are limited to 100-meter lengths, the total geography covered can be extended by cascading hubs. A break in one cable segment (other than the segment connecting the file server) may disable one network component but will not completely interrupt network operations.

3.4.2 Token Ring Standard

With major input from IBM, the IEEE 802.5 standard was developed for token ring networks. Its past popularity stems in large part from IBM's early adoption of the token ring structure as its standard for connecting mainframes, minicomputers, and PCs as peers on LANs and wide area networks (WANs). In 802.5 networks, a continuous stream of data-carrying tokens passes around the network (in one direction) from one component to the next in a circular, or ring, fashion. When one network component needs to communicate with another, it waits for an empty token to pass and inserts its data into that token. The receiving component then detects the passing token with data addressed to it and extracts that data. Throughput on token ring networks runs at either 4 Mbps or 16 Mbps.

STP cabling connects network components to a MAU. While the physical topology is similar to a star network, the ring is actually maintained as tokens are passed from one network component back to the MAU and out to the next component and then back to the MAU and so on (Figure 3.7).

Figure 3.7 — Network components: token ring STP cabling configuration.





Figure 3.8 — Network components: token ring STP cabling configuration.



The physical layout of a token ring network makes it relatively easy to add new components or move existing network components. However, the STP cable used in token rings is more unwieldy and expensive to install than UTP. UTP cabling also connects network components to a MAU, usually with a RJ-45 connector at each end of the cable (Figure 3.8). Although configured in a physical star, the failure of one segment attaching one component can disrupt the entire network since the network relies on each component's ability to receive and retransmit tokens.

3.4.3 ARCnet Architecture

ARCnet was developed in the late 1970s by Datapoint. Unlike Ethernet and token ring, ARCnet has never been officially designated as a standard by the IEEE 802 committee. However, its early acceptance, continued support, and further enhancement by a number of manufacturers have made it one of the significant industry standards for networking.

Like token ring, ARCnet utilizes tokens to notify network components when they can transmit their data packets or frames (Figure 3.9). Each ARCnet station is identified by a unique address (from 1 to 255) that is usually set on the NIC by the network installer. The station with the lowest active address on the network becomes the traffic controller, sending out tokens to each station address in turn, granting that station permission to broadcast its data. While ARCnet transmission rates are a relatively low (2.5 Mbps) the contention-free nature of its token-passing scheme enables it to perform well in heavy LAN traffic situations. Datapoint recently introduced a 20 Mbps version of ARCnet which could compete with Ethernet if manufacturers of wiring centers support it.

Figure 3.9 — Network components:
ARCnet RG-62/A coax cabling configuration.





Figure 3.10 — Network components: host terminal configuration.



3.5 ***Types of Networks***

There are three basic network types. Each is defined by where processing is performed, data is stored, and resources (e.g., printers, modems, etc.) are attached.

3.5.1 **Host-Terminal Networks**


The host-terminal network is characterized by a host computer (mainframe, mini or micro) with dumb terminals (or microcomputers emulating terminals) attached (Figure 3.10). The host handles processing for each operator on the network while the terminals provide screen output and keyboard input. Host-terminal is the oldest network type since it is the network historically used by mainframes and minicomputers.

3.5.2 **Client/Server Networks**

Client/server networks are characterized by microcomputers operating as intelligent client workstations connected to a computer operating as a

server (Figure 3.11). The server provides services (typically file and printer access but also application or communication access) to the intelligent workstations. Since the workstations perform the majority of the processing, this is considered a distributed processing network (i.e., processing is shared across a number of network components). Distributed processing generally provides more flexibility, since workstations can easily be customized to meet the specific needs of their operator and their applications. The central server can be optimized to provide common access to critical resources. Large client/server networks may consist of many servers and hundreds or thousands of workstations.

Figure 3.11 — Network components: client/server configuration.



Clients

Perform majority of processing

Server


Provides majority of file / data storage (*file server*)

Can provide shared access to printers (*print server*)

Can provide shared access to fax and modem (*communications server*)

Can provide central processing of some applications (*applications server*)

Figure 3.12 — Network components: peer-to-peer configuration.



Each computer can act as server and client - where each computer can use its own resources (e.g., programs, data, printers, modems, etc.), can also share those resources with other computers on the network (acts as a server), and can access the resources of network computers (acts as a client).

3.5.3 Peer-to-Peer Networks

Peer-to-peer networks are characterized by two or more microcomputers connected together and running network software (Figure 3.12). Each computer performs its own processing but may access or share hard disk, printer, application software, data, or resources with the other computers on the network. Like client/server networks, the distributed processing in peer-to-peer networks provides flexibility. The approach of combining client/server functions in one workstation usually works best in smaller networks where the need to share resources is not great.

3.6 Network Operating Systems

There are several network operating systems (NOSs) available from different manufacturers. Depending on the product, they may be host-terminal, client/server, or peer-to-peer systems, nearly all of which work with Ethernet/802.3, token ring, and ARCnet and their associated topologies.

3.6.1 Novell NetWare

Novell developed NetWare, the Goliath of the NOSs, which represents about 60% of the PC network market share (about 4,000,000 people on 400,000 LANs). Since its development in the early 1980s, NetWare's performance and features have been consistently enhanced. Currently, the

two main client/server network operating system product lines are NetWare 3.12 and NetWare 4.1. NetWare 3.12 is the latest version of Novell's 32-bit network operating system. It provides excellent support for file and print sharing, security features, and a wide range of hardware and software applications. NetWare 4.1 adds directory naming services, 1,000-user capacity, Windows-based administration, and data compression for increasing disk capacities. Apart from the technical merits of their products, Novell's strength is in its nearly universal acceptance by hardware and software developers. Because it represents the predominant share of the LAN market, manufacturers who hope to succeed in the LAN marketplace are highly motivated to ensure their products are compatible with Novell NetWare operating systems. Consequently, a large array of hardware and software is available for the NetWare network operating system.

3.6.2 Microsoft LAN Manager and Windows NT Server

Microsoft's (and its early collaborators) first attempts to offer a network operating system were less than successful. Microsoft has since addressed the shortcomings of its initial product. It currently markets two high-performance client/server network operating systems – LAN Manager 2.2 and the more recently released Windows NT Advanced Server 3.5. While LAN Manager 2.2 represents a significant improvement over its predecessor, Microsoft's vision for the future of a network operating system now clearly lies with Windows NT Advanced Server and its successors. The Windows NT Advanced Server supports symmetrical multiprocessing (use of multiple processors on one server), reduced instruction set computing (RISC), and Alpha and Intel processing platforms, and centralized management of users and groups across multiple, connected servers. Microsoft has the resources and appears to have the commitment to ensure that their NOS will attain a larger market share and compete with Novell.

3.6.3 Banyan Systems VINES

Banyan Systems developed VINES (VIRtual NETworking Software) as a high-end NOS with strong internetworking features. It offers a global naming service (StreetTalk), solid security features, symmetric multipro-

cessing support, and the ability to connect multiple file servers via a number of communication alternatives. It is a mature product with extensive third-party support. The system is especially worth consideration by those who need to connect widely separated servers.

3.6.4 Performance Technology POWERserve

POWERserve is Performance Technology's dedicated 32-bit NOS that provides support for up to 255 users. This NOS is fast and reasonably inexpensive compared to other comparably featured systems. It includes a global naming service, a sophisticated e-mail package, and numerous network management tools. Its cost, performance, features, and conformance to Microsoft standards and products make POWERserve worthy of consideration.

3.6.5 Digital Equipment DECnet and Pathworks

Introduced by DEC in the late 1970s, DECnet was designed to link various kinds of computers and share resources (e.g., printers, data storage devices) through their parallel interfaces. A significant drawback to the parallel interface was that it imposed a maximum distance limitation between computers of 9 meters. This was significantly extended in the 1980s when DECnet was implemented on Ethernet, the network architecture codeveloped by DEC. DECnet continues to represent a viable NOS for interconnecting DEC PCs, terminals, and minicomputers. DEC subsequently created Pathworks, a family of network products that integrate PC clients running DOS, Windows, or OS/2 with DEC's VMS, UNIX, or OS/2 server systems.

3.6.6 DOS-based LANs

Servers on DOS-based LANs can run their own applications while simultaneously making network resources (e.g., hard disk, printer) available to other workstations. These NOSs are usually easy to install and manage. However, the server PC will experience some degradation in performance due to the allocation of a portion of its random access memory (RAM) to the network operating system and the processor dividing its time between local and server operations. There are several popular DOS-based systems available.

3.6.7 UNIX

UNIX is a multitasking operating system initially developed by AT&T in the late 1960s and early 1970s to run on DEC computers that were controlling telephone switching equipment. It was subsequently adopted by many universities which, in turn, led to its wide-spread use in scientific and calculation-intensive applications. The UNIX operating system can be installed on a minicomputer or PC and can support either dumb terminals or PCs as workstations. The UNIX system, or one of its variations, has a large and loyal installed user base, and offers an efficient environment for many applications.

4.0 NETWORK APPLICATIONS

When considering a new network or planning the expansion of an existing one, a thorough review of the intended applications should be performed. The axiom in systems planning is: decide on the software first (i.e., what applications are needed, now and in the reasonable future). Only after selecting the applications and understanding their requirements can you begin to select the network operating system, hardware, cabling topology, and LAN architecture that will best meet your needs. Beyond these basic software application and hardware decisions there are a few general categories of network applications that may be worth consideration. These are discussed below.

4.1 *Resource Sharing*

A major justification for the installation of most networks is their ability to share resources among users. Sharing resources avoids expensive duplication of hardware and software. The most commonly shared resources are hard disks (and the data they contain) and printers. Increasingly, modems, computer-based fax, and CD-ROMs are also being shared among a large number of users on a network.

Software applications can also be shared on a network. Since software manufacturers generally license their products "one user per license at any given time" it is possible to save on the purchase of software by purchasing only as many licenses of a given application as are likely to be in use at a given time. Since it is unlikely, particularly in large networks, every user will be simultaneously using a word processor or spreadsheet, the total number of packages can be reduced. Finally,

data can easily be shared across a network. There are many types of data whose availability would benefit numerous users (e.g., accounting and billing data, patient data). Keeping data in one location (on a file server) facilitates centralized updates while ensuring it will be available to network users.

4.2 *Electronic Mail*

A network can significantly facilitate effective communications between its users. Electronic mail, or e-mail, is an application that enables a user to broadcast messages with attached files (e.g., text, graphics, audio) to an individual or group on the network. Used properly, e-mail can efficiently disseminate and exchange information. With proper configuration, e-mail can cross beyond the limits of local or even organization-wide networks and exchange information with individuals and organizations through public access systems like Internet.

4.3 *Group Scheduling (Resource Management)*

The scheduling of people, facilities, and equipment in most organizations is usually a hit-or-miss proposition. Scheduling programs designed for network use allow users to both determine availability and reserve the time of other people, facilities, and equipment.

4.4 *Wide Area Networking*

Wide area networks (WANs) link multiple local area networks (LANs), thereby giving users access to information and resources across the combined system. Properly configured, WANs permit hospitals, physicians' offices, satellite clinics, health maintenance organizations (HMOs), or any healthcare organization to enter or review patient records, laboratory results, medical images, and other information from any point in the LAN. The means of interconnecting LANs into a WAN are discussed in Sections 4.4.1 and 4.4.2.

4.4.1 *Scheduled or On-Demand Connections*

LANs that do not need constant contact with each other (i.e., users from one LAN do not need to have real-time access to information or resources on another) can maintain intermittent contact. Typically, one

LAN will access another via telephone connections, which can be established on a scheduled or as-needed basis. Once a connection is made, the LANs exchange information. The intermittent nature of these connections and the relatively low transmission speeds attainable through most telephone lines make this approach most suitable for e-mail and file updates between LANs.

4.4.2 Continuous, Dedicated Connections

Some applications demand continuous high-speed links between LANs. These links facilitate real-time accessing and processing of data across the WAN (regardless of which LAN the programs and data actually reside in) and rapid transmission of large amounts of data (e.g., high-resolution graphics, video, multimedia). Continuous and high-speed links come at a premium. The interconnection services may be available through an independent service, but most often they are obtained through local, regional, or long distance telephone companies.

4.4.2.1 T-1 Lines

T-1 is the designation for a connection rated at 1.44 Mbps. These connections traverse hundreds or thousands of meters over leased long distance facilities. Fractional T-1 service is also available when full 1.44 Mbps service is not required.

4.4.2.2 Integrated Services Digital Network

Available now in limited areas, all regional phone companies are committed to the eventual implementation of Integrated Services Digital Network (ISDN) telephone service, which will provide 64 Kbps to 128 Kbps transmission.

4.4.2.3 Metropolitan Area Networks

Still being planned by local phone companies or under development in metropolitan areas, metropolitan area networks (MANs) cable systems (typically fiber-optic) will provide high-speed, low-noise services to connect LANs within their coverage area.

4.4.2.4 Fiber Distributed Data Interface

Fiber distributed data interface (FDDI) is a fiber-based system suitable for intracity networks. FDDI services will be built by some organizations with an appropriate level of need, and excess capacity may be offered to others. FDDI will operate with a sustained data throughput of 100 Mbps.

5.0 NETWORK HARDWARE

Most networks contain several key hardware components. The components described in the following sections are typically found in Novell NetWare, Microsoft NT, or DOS-based local area networks (LANs).

5.1 Servers

The server is a computer configured to provide centralized services for users on the network. Typically, these services are shared file and printer access, but they can also include shared access to fax/modems or to database engines. One computer can be set up to act as a combination of servers (e.g., both file and print server) on one network, or one network may have multiple servers, each specializing in a different service (Figures 5.1a and 5.1b). The right configuration of single, multipurpose or multiple, single-purpose servers depends on network size and utilization. Small, low-traffic networks will probably work well with a single server providing multiple functions. Larger enterprise networks will perform best with multiple servers dedicated to various functions.

5.1.1 File Server

The most basic server is a file server, which stores files in a central location and sends them as requested over the network to a workstation. It also creates, modifies, and deletes files on the server as requested by a workstation.

While the specific configuration for a server depends on its function, a good multipurpose server typically has an Intel 80486 or Pentium processor, a large hard drive (preferably 1 gigabyte or larger), 16 megabytes or more of RAM, a fast data bus (e.g., PCI), several expansion slots (four to six), a network interface adapter designed for the computer's fast bus, and a fairly large power supply (e.g., 300 watts). Servers with multiple

processors have recently become available the popularity of which will increase as the number of network operating systems and applications that can take advantage of symmetrical processing.

5.1.2 Print Server

The second most common type of server is the print server. It manages requests from workstations for print services, queuing, and routing print jobs to available printers on the network. A print server permits a group of users to install and access a variety of printer types (dot matrix, lasers, ink jet, color) while eliminating the need to attach one printer to each user's computer.

Figure 5.1a — Network components with a single multipurpose server.

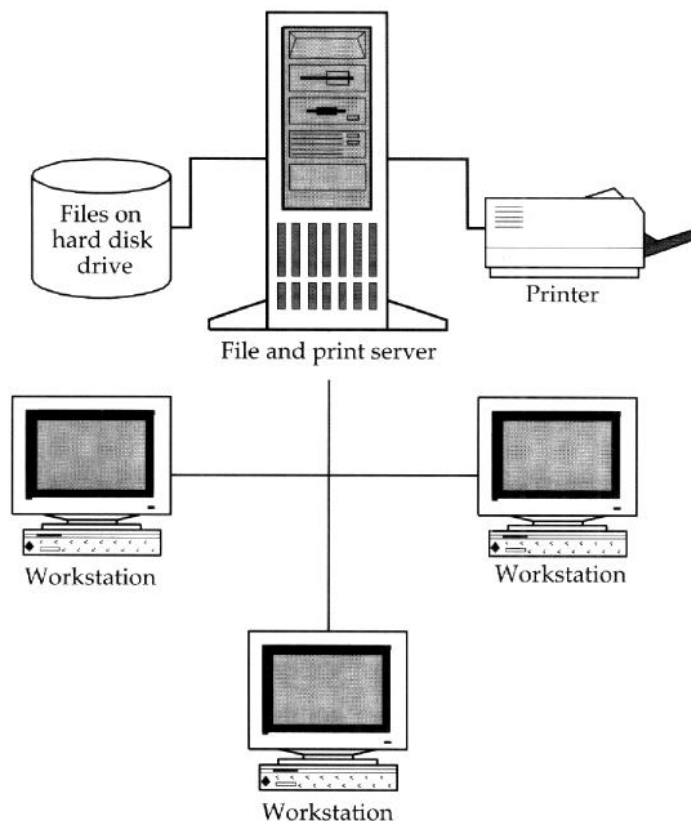



Figure 5.1b — Network components with multiple single-purpose servers.



5.1.3 Communications Server

A communications server provides workstations with shared access to one or more modems or fax/modems. As with the print server, a communications server can minimize costs since fax/modems and phone lines do not have to be installed at each workstation in order to give their

users access to these services. Communications servers configured as fax servers enable network users to dial out, and they receive and route incoming fax calls to a specific user. Another important feature of communication servers is their ability to provide complete network access to users who are located off-site. Remote users can dial in to modems on communication servers and can run any program on the network and access any data as though they were accessing the network from an on-site workstation.

5.1.4 Application Server

Application servers are another type of server whose popularity is increasing as networks are used in more critical operations. Application servers do most of the processing associated with an application program on the server. For example, an application server where the application is a database engine would perform the data searches and report generation on the server. The workstation would run front-end software, passing along requests for searches and reports to the database engine on the server.

5.2 Workstations

The workstation in host / terminal networks provides only keyboard input to and screen output from the host. Processing occurs on the host, not the terminal, side of this network. Client / server and peer-to-peer networks make use of intelligent workstations, where the workstations perform the majority of the computer processing. While the configuration of these workstations will depend on their intended use, most will have:

- A 32-bit processor (50 MHz minimum).
- Sufficient RAM (8 Mbytes minimum) to run programs and, a multi-tasking environment such as Microsoft Windows 95.
- A small hard disk (200 Mbytes) for programs to store temporary working and configuration files (most other program and data files are stored on the more capable file server).
- A sufficient number of expansion slots and bays to handle peripherals (e.g., scanners, fax / modems).
- A network interface adapter.

5.3 Backup Systems

Making regular backups of programs and data is a critical function on a network because computers are susceptible to a variety of failures. The hard disks that store both programs and data are mechanical devices and, as such, are particularly vulnerable. It may be impossible to recover any information following a catastrophic hard disk failure. The implication of this kind of failure is serious enough when it involves one user's PC. It can be devastating to a business when it is the file server's disk that contains critical business records and the programs and work of many users. Magnetic tape is the media used almost universally to store backup information; drives and tapes are available in sizes ranging from 100 Mbytes to many gigabytes. Network backup systems may either be server or workstation-based. Some server-based backup systems can backup workstations as well as the file server. Regardless of the type of backup system, it is important that it provide:

- Automatic or easy backup. The more manual intervention a backup system requires, the less likely it is that it will get adequate use.
- A set of multiple backups. Backing up to a different tape each day of the week provides an additional margin of safety should one tape be bad or should several backups get made before program / data problems are discovered. Cycling one backup through off-site storage provides some safety against theft or facility damage.

5.4 Network Security

Since networks are designed to give multiple users access to information that may be critical and sensitive, network security is usually a significant consideration. Nearly all network operating systems (NOSs) feature security systems that work over and above any security systems that individual applications may employ. The most basic of these systems limit user access, based on the user's name and password, to the specific program and data areas on a file server. Network users may also be given read-only rights to certain files and they may have date and time-of-day restrictions placed on their access. Network security systems can restrict access to confidential information and can restrict user access to minimize the likelihood of accidental or intentional damage to programs and data.

5.5 *Fault-Tolerant Systems*

Networks are usually critical to some degree due to the nature of their use in business, security, or safety applications. To improve their reliability and survivability, they usually incorporate some degree of fault tolerance. Fault tolerance in this context means anticipating failures of key critical network components and providing a means of rapid recovery from those failures. A network designed with these considerations, therefore, becomes fault tolerant and can tolerate a fault without any significant loss in functionality.

5.5.1 Uninterruptible Power Supplies

Power interruptions are a common source of network problems. Interrupting power to a file server, hub / concentrator, or other critical network component can disrupt the entire network, including the users and their operations. The maintenance of power is important because a sudden interruption can corrupt data files. The uninterruptible power supply (UPS) is designed to temporarily provide continuous power to critical network components in the event of a power loss or brown-out. When a power problem occurs, the UPS senses a change and switches over to an inverter via a battery. Many UPSs have the ability to communicate with the network operating system which, in turn, can notify users of the change in power status. While battery operating time may be limited to a few minutes, it is usually enough to enable users to finish operations and log off properly. Those UPSs with the ability to communicate with the network operating system can, after a predetermined period (and before the battery runs down), execute a proper shutdown of the server.

5.5.2 Redundant Systems

Critical networks that cannot tolerate interruption can utilize redundant components to ensure that a failure of any one key component will not disrupt network operations. Redundant systems can include hard disks, power supplies, and even entire duplicate servers and wiring centers. To avoid interruptions, these redundant systems may provide for automatic switching from a defective component to a redundant one when a failure occurs. Alternatively, many of these components can be “hot swapped” when they become defective (i.e., the defective component can be replaced while the network continues to operate). One of the more commonly implemented redundancies involves hard disk subsystems. There are five defined levels of Redundant Arrays of Inexpensive Disks (RAID) that provide some hard disk subsystem redundancy.

5.5.2.1 RAID 1

RAID 1 involves the use of disk mirroring. Disk mirroring uses a primary and secondary disk drive (preferably identical models) on one controller. All data written to the primary drive is also copied (mirrored) to the secondary drive. A failure of the primary drive would cause the secondary drive to immediately take over with no loss of data. A variation on disk mirroring is called disk duplexing, which uses a primary and secondary disk drive and controller. All data passes through both controllers and is written to each disk drive. Under normal operations, system performance is enhanced through simultaneous read / writes. When either a controller or disk fails the second controller / drive combination can take over.

5.5.2.2 RAID 2

In RAID 2, bits or blocks of data destined for storage are divided and written to several disk drives in parallel streams. Some number of disk drives out of the total group are reserved to store error-correction information regarding the data written to other hard disks. Should any one drive fail, it can be replaced with a new drive and its data reconstructed from the data and error-correction information stored on the remaining drives.

5.5.2.3 RAID 3

RAID 3 is similar to RAID 2, but stores parity checking rather than error correction information on some of the drives in the array. Restoring data on a replacement drive takes more time when using parity checking rather than error correction information, but the parity checking data requires less drive space and, therefore, can mean fewer drives are needed in the array.

5.5.2.4 RAID 4

RAID 4 stores data on multiple drives serially on a sector-by-sector basis (rather than in bits and blocks as in RAID 2). A single drive in the array is dedicated to the storage of parity information.

5.5.2.5 RAID 5

RAID 5 is similar to RAID 4 but eliminates the dedicated parity drive. Rather than place parity information on a single dedicated drive, it rotates parity information across all drives in the array. RAID 5 is a popular implementation of RAID technology because it provides fault tolerance and redundancy with as few as two drives.

5.6 Bridges

Bridges connect similar or identical LANs (such as Ethernet-to-Ethernet). These connections occur at the Media Access Control (MAC) sublayer (within the data link layer) of the Open Systems Interconnection (OSI) model. Bridges can act as effective filters of LAN traffic since they only forward traffic across the bridge when it is addressed to network components located on the other side. A bridge can be a dedicated computer or a piece of hardware designed for the purpose. The bridge uses a fast processor and software to examine the destination address of data moving on one LAN segment, passing only information to the next LAN segment when addressed to a component on that segment.

5.7 *Routers*

Routers are more complex linking devices than bridges. They have the ability to examine and direct traffic (by adding information to each data packet) and can link dissimilar LANs (such as ARCnet, Ethernet, and token ring). Routers connect LANs together at the network layer (layer 3) of the OSI model. Unlike bridges, routers are not transparent to stations on the LAN — stations on one LAN segment must specifically address packets or frames to a router for transmission to an address on the other side.

5.8 *Gateways*

Gateways, which function at the high end of the OSI model (i.e., applications, layer 7), repackage or convert data sent from one network to conform with the format of the application on the receiving network. For example, an e-mail gateway would take a message from one network and translate its components (recipient field, cc: indicator, notification of receipt, priority level, etc.) to the format of the e-mail package on the second network.

6.0 NETWORK MANAGEMENT

The major categories of network administration are considered to be the following:

- Configuration management (device inventory, changes to the configuration, network reconfiguration, etc.).
- Fault management (device availability, device control—bringing devices up or down).
- Capacity management (measuring and changing devices throughout).
- Usage accounting.
- Security management.

The management of a network, particularly an enterprise-wide internetwork, is a complex set of activities. Networks are dynamic applications and users are continually being added and deleted. Both hardware and software get upgraded. Network components can fail, requiring repair or replacement. Network segments become overloaded,

lose efficiency, and need to be reconfigured. For purposes of software licensing, security, cost allocations, and planning it is important to understand how the network is being used by end users. Inventory control is also an important issue; the performance of a network can be greatly impacted by the hardware and software on the individual workstations. Changing the software version or hardware components can prevent use of the network.

6.1 *Single Vendor vs. Multivendor Environments*

Many network components have become standardized to the degree that products from different vendors generally work well together and can be mixed and matched in a variety of combinations. Going to a single source or vendor for most networking needs should ensure that the network products integrate well together. On the other hand, relying on a single vendor typically means that compromises are made since one vendor's products do not usually represent the best in every category. A single-vendor approach is the safest in those situations where network administrators are relatively inexperienced and nontechnical. However, when a single vendor is selected, the vendor's history and references should be thoroughly examined. More experienced and technically-oriented network administrators may be comfortable dealing with a variety of vendors so they can select the vendor offering the best combination of performance, support, and price.

6.2 *Consultants*

The growing complexity of networks is causing companies to turn to consultants and other third parties for assistance in network design, implementation, and management. In some cases this leads to complete outsourcing of the network. Outsourcing is increasingly feasible as network standards spread, causing networks to look and behave more like a utility. It is this utility-like aspect of networks that makes them attractive to the telephone companies who would like to provide more local area network (LAN) and wide area network (WAN) services to customers.

6.3 *Components to be Managed*

Network entities that need to be managed can be divided into several classes:

- Devices – hubs, routers, bridges, gateways, and servers.
- Segments – cable segments.
- Applications – software employed by end users.
- Workstations – the individual hardware components within a workstation – network interface card (NIC), central processing unit (CPU), ports, hard drive, pointing devices, etc.).

6.3.1 Device Management

It is important to know the status of hubs, routers, servers, gateways, etc. This can be done by placing software (or firmware) in the device, storing the parameters in the device, and communicating this data to a centralized location for management. When tracking the status of devices, two issues arise – frequency of status reporting and medium of status reporting. In the former case, it is important to be able to specify the frequency with which the device reports its status – too frequent reporting consumes scarce bandwidth. In the latter case, it can be useful to do status reporting over a different medium (out-of-band reporting) – if a device broadcasts its status over the same network as the end-user applications (in-band reporting) a network problem may prevent the device from transmitting its status. Frequently, something as simple as a RS232 connection is used for out-of-band reporting.

6.3.2 Simple Network Management Protocol

The simple network management protocol (SNMP) was developed for the transmission control protocol / internet protocol (TCP / IP) world in the late 1980s. SNMP was designed to manage the multivendor networks that comprise the Internet. It drew heavily upon the network concepts of open system interconnection (OSI), which had been defined but not implemented. It is a layer 7 protocol developed for TCP / IP, which has been ported to run over other transport and network protocols. SNMP has experienced rapid growth during the first half of the 1990s, as companies try to attain control of their expanding networks. The basic objective of SNMP is to monitor and control network devices.

The architecture of SNMP is based on the concept of agents and management information bus. Each network device (hub, concentrator, router, etc.) is capable of monitoring the packets of data that pass through it, storing this data in a management information bus (which is a database of the characteristics of this device). The data includes device setting, status, timer, topology, throughput, usage, etc. A program called "the agent" runs in the device, collecting data and updating the device's management information bus at specified intervals. For its physical implementation this management information bus can use any of several different technologies: relational, object-oriented, flat files, or RAM; and the agent may be implemented as either software or firmware. While at least 60% of all devices sold are SNMP-compatible, not all are. These non-SNMP devices can be controlled through SNMP's proxy agents. A proxy agent functions as a gateway to a non-SNMP device, translating SNMP commands into proprietary commands.

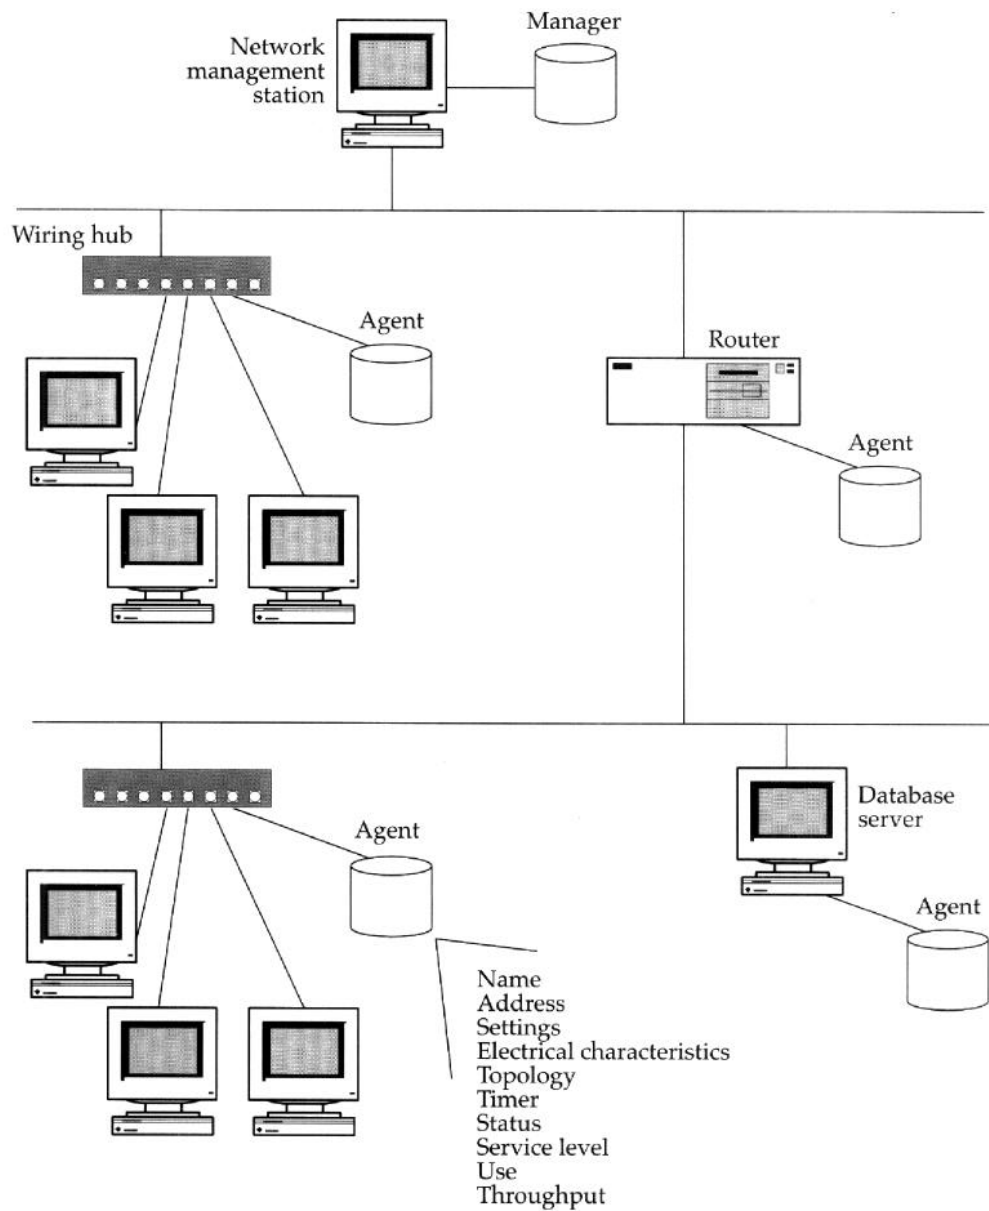
A management workstation, which has a program called "the manager," sends both queries and commands to the agent of each device. The queries ask about device settings and status. The commands direct the agent to change settings, bring the device up or down, etc. (see Figure 6.1).

One weakness of SNMP is its lack of security, and so a second generation, SNMP2, is being developed. In addition to improved security, SNMP2 allows vendor extensions to the management information bus, which permits individual vendors to provide features without interfering with SNMP.

6.3.3 OSI CMIP

The OSI equivalent of SNMP is the common management information protocol (CMIP). It has greater functionality but, like all OSI protocols, is less efficient and less widely implemented. SNMP will eventually be replaced by CMIP for the TCP/IP suite.

Figure 6.1 — Architecture of SNMP.



6.4 Vendor Solutions

Realizing that the trend is increasingly toward multivendor networks, vendors of networking products are developing new products for the management of this complex environment. These products are usually based on SNMP. The vendor provides added features for the network manager: centralized / distributed capability, windowed workstation, graphic reports and diagrams, statistical analysis, and other planning tools. A major issue arises when a hospital acquires network products from multiple vendors and each vendor supplies their own network management workstation and software. The approach in this situation, which is becoming increasingly common, is to acquire a manager of managers, a program that sits above the other vendor-specific managers and integrates them.

6.4.1 Netview

IBM's Netview is designed to manage both systems network architecture (SNA) and non-SNA networks together. It permits translation between them so that they can be managed from a centralized workstation (a focal point in IBM's terminology).

6.4.2 Open View

Hewlett-Packard's Open View has a distributed architecture and is designed for both LAN and WAN environments. It employs a graphical windows user interface.

6.4.3 SUN Net Manager

SUN's Net Manager runs on a UNIX workstation and can be either centralized or distributed.

6.4.4 Other Managers

DEC was actively engaged in producing its own network management architecture, enterprise management architecture (EMA), but abandoned its efforts in 1994. Distributed management environment (DME) was developed by OSF; however, lack of support has caused cancellation of this project.

6.5 Cable Management

In addition to knowing the status of network devices, it is also important to know the location and condition of the cable runs that constitute the network. There is a wide variety of diagnostic products, including cable scanners and protocol analyzers, used for network troubleshooting.

6.5.1 Cable Scanners

The cable scanners for coaxial and twisted-pair cable can locate cable breaks and faults (within 0.3 to 0.6 meter), signal degradation, crosstalk, and various other cable problems. In 1994, cable scanners cost between \$1,000 and \$8,000.

6.5.2 Protocol Analyzers

Protocol analyzers are diagnostic tools used by network troubleshooting experts to analyze the data packets exchanged between network components. These analyzers work with Ethernet, token ring, and ARCnet networks, and the various network protocols (IPX, IP, XNS, etc.). In 1994, protocol analyzers consisted of both hardware and software and cost anywhere from \$1,200 to \$25,000.

6.5.3 Cable Management System

A cable management system (CMS) is a computer-aided design (CAD) package that allows one to track the cable runs, showing all the connections along the way. This helps to identify unused cable and available ports.

6.6 Application Management

For several reasons – licensing, chargebacks, and security – it is important to know who is using what software on the network.

6.6.1 Application Metering Software

Metering software is LAN-based software that monitors the use of applications (e.g., word processing, spreadsheet, database, etc.) by network workstations and provides a report of usage history. Metering software can also limit users on an application to a number that has been preset by the network administrator. Network administrators can thereby sig-

nificantly reduce their software budgets by purchasing only the number of licenses that would be in use at one time. This software can also do usage accounting (bill users for their use of network resources). Trying to identify the precise amount of usage of each network component at the level of the individual user, or even the department, requires much effort. The hardware and labor in many cases would cost more than the benefits derived from the accounting system. Therefore, many companies treat the network as a utility, an item of corporate overhead, and only charge users for the devices in their immediate area.

6.6.2 Security and Viruses

If a network is the unfortunate target of a virus, lack of preparation can be devastating. Restoring from backups after a virus flare-up may be ineffective if the virus was present and dormant during the backup process. The best protection against viruses is the use of virus monitoring software coupled with an extended set of backups. Good virus-monitoring software will run continuously in the background, examining all files as they are being written to the network. The antivirus software will block any recognized virus from being written to the network and notify the administration when it is detected. It is also a good procedure to ban the use or installation of unauthorized or personal software on the network.

Network break-ins (incursions) can occur at the level of hardware or software. The cables can be tapped (difficult to do with fiber, however), or the electromagnetic leakage can be read by special monitoring devices. On the software side, the very purpose of internetworks and enterprise networks is to provide connectivity and access to a wide range of users, and this makes them susceptible to break-ins. Techniques to prevent incursions include encryption of messages and names, dialing back to the person requesting access (instead of passing them through), and placing passwords on operating systems, file directories, applications, and even modules.

6.7 Workstation Inventory Management

A change to a workstation's configuration (either hardware or software) may cause it to stop working with the network. For this reason it is important to control the workstation inventory. Several vendors have developed software that consists of two parts – a program that resides on a

server and one that resides on the workstation (a terminate and stay resident for DOS-based machines, and a remote procedure call (RPC) for other operating systems). Periodically, the manager program will poll the individual workstations, recording their inventory and comparing it to the database to determine whether the workstation's configuration has changed. This information is then available for troubleshooting and planning.

7.0 EXTENDED NETWORKS

The concept of an extended network is much broader than that of a wide area network (WAN). Many hospital chains have employed WANs for the periodic transmission of performance data to corporate headquarters or to connect geographically dispersed terminals to a centralized processor. In the past, there has been some data exchange between individual hospitals and suppliers of medical supplies (e.g., the Baxter ASAP system), between hospitals and insurance companies, and between hospitals and the federal government. A recent and major trend has been the expansion of networks beyond the environments of a single legal entity, in a systematic way, community healthcare information network (CHIN). In addition, communications companies, both local and long distance, are beginning to provide networking services beyond the simple transmission of a message.

7.1 *Technology for Extended Networks*

On a local level, the way to extend a local area network (LAN) is through the use of bridges and routers. In general, bridges connect similar networks (e.g., Ethernet-to-Ethernet), while routers connect dissimilar networks (e.g., Ethernet-to-token ring). Historically, if a hospital wanted to extend its LAN over a broader geographical area, connecting to different entities, it would use the networks of the long distance carriers, a combination of leased lines and X.25 packet switching. This did not provide much throughput or flexibility. New technologies have appeared in the past 5 years to allow much greater throughput, flexibility, and functionality, making it quite feasible to achieve tighter interaction among organizations. These technologies include digital long distance lines (T-1 and T-3); improved layer 2 protocols (ISDN, Frame Relay, and the emerging ATM);

improved layer 7 protocols (e-mail, file exchange, and database access), allowing connectivity and interoperability of applications; and enterprise-wide and interorganizational routing protocols.

7.2 Community Healthcare Information Network

Historically, wide area networking was used for two major applications: 1) tying together terminals in a transaction processing system to a centralized processor (e.g., airline reservations, bank terminals, hospital admitting); and 2) batch file transfer between systems. WANs were not well suited for such things as database access across multiple systems, universal mail, etc., but this is changing dramatically (see Section 7.1).

In addition to technology, trends within the healthcare industry are driving the movement toward CHINs, primarily the attempt to get healthcare costs under control. This requires consolidation and coordination of services, fast response time, synchronization and/or integration of data, and universal access to data. This consolidation and closer cooperation among healthcare entities is seen in a CHIN. In a given community, hospitals, physician offices, reference labs, insurers, employers, medical schools, and government agencies are tied together. This can also be done at multiple levels: local, county, state, regional, and national.

It is useful to review what was said about the various layers of the OSI model. In order to achieve a CHIN, one needs to have connectivity and interoperability at three different levels: layers 1 and 2 (for basic connectivity); layer 3 (for access to network resources); and layers 4 to 7 (the application layers - for interaction with network resources). A critical component is standards for data interchange, which is discussed in Section 8.0.

The administrative aspects of CHINs are daunting. Someone will have to construct a logical model for how it all fits together. This data will be of different modalities: unstructured text, structured text, database, video, images (CAT, MRI, PET, x-ray), audio (physician comments), and waveforms. It is not out of the range of possibility that an insurer would request the DX, notes, images, and charges for a patient whose bill is in question. Future networks will have to accommodate this data complexity. It appears as though ATM technology will provide the bandwidth to do this, although it could also be done with parallel networks.

Professionals will have to devise policies on data access and use. Each of the participants in the CHIN will contribute data and use data. Usage accounting is more complex by an order of magnitude.

7.3 Public Access Systems

There are two trends among the telecommunications carriers: 1) different media are becoming consolidated – long lines, cellular, local exchanges, and cable are merging (e.g., AT&T's acquisition of McCaw Communications, a cellular vendor) and; 2) service is becoming vertically integrated (MCI is going into local markets).

In addition, vendors are providing more network services. Traditionally, third-party carriers provided multiplexing, routing, and packet exchange. Now, they are providing office automation (AT&T and MCI), and LAN services in addition to WAN services. They are turning it into a Centrex for LANs - "a LANtrex." In addition, there are many vendors like GEICO which provide gateways and clearing houses for electronic data interchange among different business entities.

8.0 HEALTHCARE STANDARDS

Historically, there has been a distinction between intramural data exchange and extramural data exchange. Extramural data exchange was performed between a hospital and its suppliers and its payers (insurance companies and government). Now, with the growth of community healthcare information networks (CHINs), which need to exchange all types of data, the distinction between intramural and extramural is becoming less clear and less relevant.

In the past, vendors like McKesson and Baxter had proprietary standards for exchange of purchase orders with hospitals. This is giving way to X.12, an American National Standards Institute (ANSI) standard for purchase orders, acknowledgments, invoices, and remittance advice.

There are several standards for the exchange of patient and clinical data: Health Level 7 (HL7), Medical Data Exchange Committee (MEDIX), American Standards for Testing and Materials (ASTM), American College of Radiology / National Electrical Manufacturing Association (ACR/NEMA), and Medical Information Bus (MIB). While there is some overlap, each of the standards deals with a different aspect of healthcare data. HL7 deals with ADT, orders, results, and billing in an

acute care hospital setting. MEDIX deals with a broader set of healthcare data. ASTM deals with laboratory data. ACR/NEMA deals with imaging and radiology data. MIB deals with a patient's physiological data. Fields that they all have in common would include: patient ID, physician ID, facility ID, etc.

8.1 **HL7**

HL7 stands for Health Level 7, referring to layer 7 of the OSI model, in which one can find the abstract specification of the message types required for data interchange. The HL7 Working Group was established in 1987. It consists of both users and vendors, and its goal is to develop standards governing the exchange of key data sets among healthcare applications. The latest version of HL7 is 2.2, published in 1994. In 1994, HL7 was recognized as an ANSI standard.

HL7 is primarily concerned with data exchange, not system connectivity and interoperability, with maintaining compatibility with existing systems. The focus is on acute care hospitals, and in particular, patient admission, discharge, and transfer (ADT), and registration data; order data, accounting data; and clinical data such as lab results. It also deals with queries and the synchronization of common master files. HL7 specifies the interfaces for systems that send or receive ADT data, orders, results, and bills. This eliminates the need for custom, vendor-specific interfaces.

HL7 does not employ modeling techniques for either events or data. Nevertheless, its model is one of events and messages. A trigger event occurs, which causes the transmission of one or more messages to all the systems that have a stake in that data. For example, all systems contain patient demographic data (medical record number, patient name, account number, etc.). When this changes, they all need to know about it (i.e., the databases have to be synchronized). A message can have one or more segments (e.g., a header record, a detail record, and a trailer record). Each segment can be divided into data elements (fields). Different types of trigger events and messages are specified for each functional area. ADT events include Admit a Patient (A01), Transfer a Patient (A02), Discharge a Patient (A03), Register a Patient (A04), Pre-admit a Patient (A05), through Unlink Patient Information (A37).

This will be illustrated with an example from ADT – patient admission. An admission event triggers the transmission of a message from the originating system to the stakeholding systems – lab, radiology,

pharmacy, billing, etc. HL7 does not try to define which system is the authoritative source. In this case, the registration system is probably the authoritative source, but HL7 does not require that it be. Nor does it care whether the systems are centralized or distributed. You could have several applications running on the same computer, or they could all be on different computers. The resulting message contains a patient ID segment, together with a header and a trailer segment, and an optional segment. Figure 8.1 illustrates this situation. A patient is admitted through the patient registration system. This event triggers a message that is sent to the lab, pharmacy, and radiology systems. The message consists of different segments (records). Some of them are obligatory (header, event type, and patient identification). Others are optional (next of kin). Many of the optional segments may repeat (e.g., observation / result). The destination systems can use as much or as little as they want. Each segment is further divided into data elements. HL7 specifies the abstract message, the syntax of the segments, the data types of the elements, and the possible value for the elements.

HL7 specifies acknowledgment messages. It also provides for batch transmission of messages. The messages may be sent in real time, or held until the end of the day. The target systems could be on different types of networks; it does not require a particular type of physical network. HL7 does not propose a suite of protocols below layer 7 (actually, its syntax notation and specification of data types are at layer 6); this is done by MEDIX (see Section 8.2).

8.2 MEDIX

MEDIX is the informal name for Institute of Electronic and Electrical Engineers (IEEE) P1157 – the committee on Medical Data Interchange. It was sanctioned by the IEEE in 1987. The goal of MEDIX is to define standards for the interoperability of heterogeneous healthcare information systems, in contrast to HL7, which focuses on the narrower issue of interfaces. Therefore, MEDIX specifies the entire OSI stack, including remote operations service element (ROSE), while HL7 deals only with layer 7. MEDIX and HL7 differ in several other ways. MEDIX is more academic and theoretical, while the HL7 committee is more practical and implementation oriented, with significant vendor participation. MEDIX uses abstract syntax notation and object modeling, including both a transaction model and a data model. HL7 avoids modeling. When comparing the two, one must conclude that MEDIX's strength is

HL7's weakness, and vice versa. MEDIX is rigorous and aims for a complete OSI protocol stack, but its implementation is dragging. HL7 has an incomplete and nonrigorous stack, with a bias to accommodating outmoded technologies, but it has been able to implement several successive versions. MEDIX and HL7 plan to merge eventually at layer 7, the abstract message definition level. A convergence committee has been established for this purpose.

Figure 8.1— HL7 triggered by admission event.


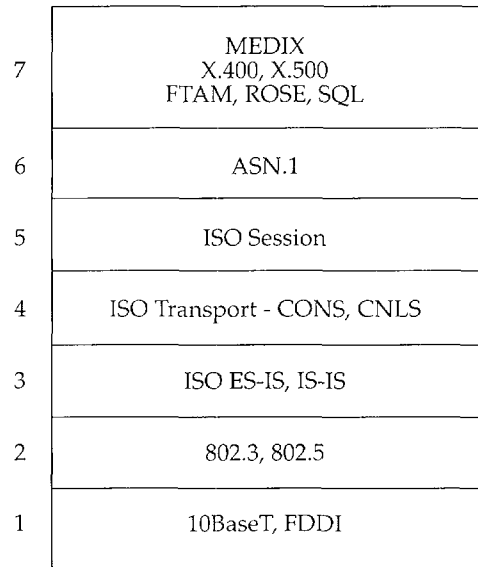


Figure 8.2— MEDIX protocol stack.



Where possible, MEDIX uses already existing OSI protocols. It intends to assemble them into an international standards profile (ISP) – a complete, seven-layer stack of OSI protocols used to provide connectivity for a given industry or function. Examples of industry-specific ISPs are manufacturing automation protocol (MAP) and technical office protocol (TOP). Figure 8.2 illustrates the MEDIX stack.

The slow adoption of OSI standards for levels 3, 4, 5, 6, and part of 7 presents problems for the MEDIX stack. MEDIX could consider substituting transmission control protocol/ internet protocol (TCP/ IP) for layers 3 and 4. However, layers 5 and 6 would still remain a problem. While layers 5 and 6 are being addressed by the middleware vendors, there is no vendor implementation of middleware that could be adopted as a standard. See Section 8.6.

8.3 **ASTM**

The Association for Testing Machinery (ASTM) has published its own specifications for the exchange of results produced by laboratory equipment (ASTM E3.11).

8.4 **ACR/NEMA**

In 1982, a committee was formed by the ACR and the NEMA to develop a standard for connecting digital imaging devices. Over time, the scope of this standard has evolved from a low-level specification of point-to-point hardware interfaces to a complete set of messages running over an OSI-standard network. These messages contain fields for control data, fields concerning the patient, and fields with the associated digitized images – CAT, MRI, ultrasound, computer radiography, etc.

8.5 **MIB**

The MIB (IEEE Standards Project 1073) is a standard for connecting medical devices to a host computer. In addition to providing the exchange of patient physiological data for storage in an all-electronic medical record, the MIB also defines standards for the remote adjustment of devices and alarms. The MIB accommodates a wide range of devices such as ventilators, infusion pumps, pulse oximeters, monitors, thermometers, and transcutaneous gas monitors.

As with MEDIX, the adoption of MIB has been slow. In the meantime, several vendors have developed products with a similar architecture, but without the full suite of OSI protocols, that is, a universal controller for all devices and a gateway to a clinical information system (CIS).

8.6 **Vendor Implementations**

Getting applications to exchange the right data, in the right way, at the right time is very difficult. There is a huge gap between the definitions of abstract message types and the protocol stacks that must implement them. Not all vendors have adapted their systems to accept and transmit transactions in standard format. Moreover, few hospitals have the expertise to install the software, to say nothing of developing it. Once again, issues of middleware greatly complicate the effort (see Section 2.8.1).

Seeing an opportunity, several third-party vendors have created gateways linking diverse applications. These systems can accept a transaction, and specify in a routing table those systems that should receive this message. They can translate from one protocol to another (IPX to IP) and from one medium to another (RS232 to RG58). Since such systems are clearing houses and are table-driven, they can be used to copy transactions for a consolidated, decision support database. This is useful in

cases where the source system cannot support a broadcast of an HL7 transaction. The source system only has to send one transaction to the store-and-forward box, which takes care of broadcasting, routing, and acknowledgements.

9.0 CASE STUDIES

9.1 *M. D. Anderson Cancer Center*

The University of Texas M. D. Anderson Cancer Center is a tertiary-care facility located in Houston, Texas. It is part of the University of Texas system and is dedicated to research and the treatment and prevention of all forms of cancer. In 1993, it had 586,979 outpatient visits, 18,701 admissions, and it registered 13,539 new patients.

M. D. Anderson receives only about 13% of its budget from the state of Texas, most of which is earmarked for indigent care. The remainder of its operating expenses must be recovered through the provision of services to paying patients. Like other healthcare institutions, M. D. Anderson has been affected by the changing competitive and political environment and it is actively seeking HMO business. It has also started to extend its presence, setting up a remote clinic (with teleradiology) in McAllen, Texas, and a full-service hospital with its name in Orlando, Florida. In the near future it will have a facility in the Fort Worth area. On the cost side, it had to eliminate several hundred positions. Related to these two areas of cost and service, it has launched a large-scale "business process engineering" project to identify areas where operational efficiency can be increased. In addition, there is increased demand for integrated financial and clinical data to aid in analysis of cost, quality, pricing, and marketing. All this has enormous implications for the networks at M. D. Anderson, in terms of integration and extending connections to more entities.

9.1.1 Overview of Networking

In 1986, M. D. Anderson installed its first Ethernet network. It remained a bridged network until 1993, at which point routers (to keep local traffic off the Ethernet backbone) were installed. M. D. Anderson installed

structured premises wiring starting in 1989. All existing buildings were retrofitted for this, and all new buildings have this in their design. Wiring hubs are located in closets, and fiber cable is used in the vertical risers. Currently, they are examining the feasibility of installing router cards for intelligent hubs.

Historically, M. D. Anderson was a large user of IBM mainframes and systems network architecture (SNA). However, the number of SNA devices has been decreasing from 3,000 SNA terminals in 1992 to 1,800 in 1994. In addition to SNA, a mixture of protocols are used: DECnet, TCP/IP, IPX, and Appletalk. Overall, M. D. Anderson is moving to TCP/IP and eventually to open systems interconnection (OSI), but will retain SNA. In order to simplify network management, M. D. Anderson would like to phase out DECnet and Appletalk, keeping just TCP/IP and IPX.

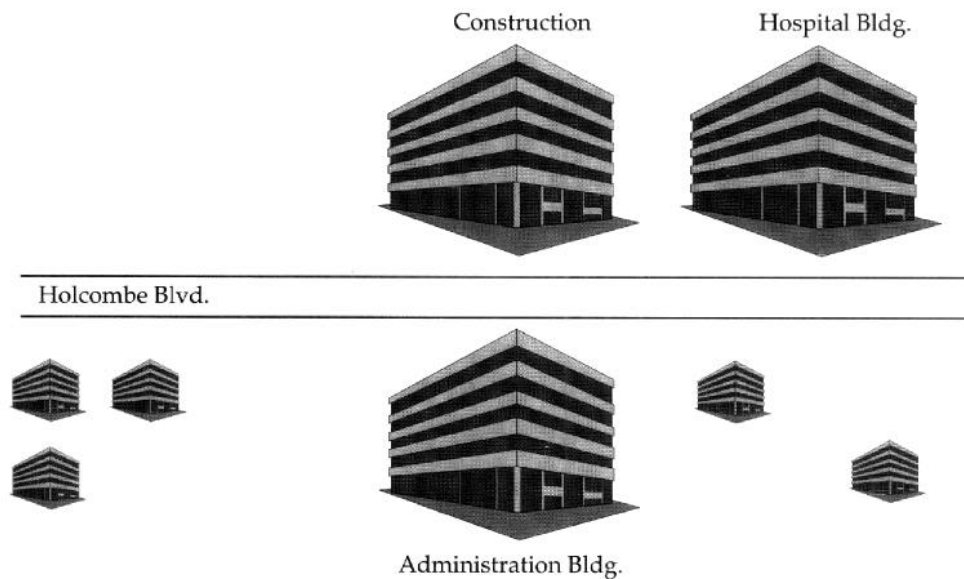
M. D. Anderson has 87 separate Novell networks (using IPX) directly connected to the Ethernet backbone. IPX and TCP/IP packets coexist on this backbone. As stated, they are installing routers to isolate local departmental traffic. When M. D. Anderson communicates with the state government in Austin it tunnels the IPX in IP packets.

There are several different gateways between local systems and the SNA hosts. The Ethernet backbone contains several system application architecture (SAA) gateways, connecting PCs to the IBM mainframe. There is also an interlink gateway connecting DEC systems to the mainframe.

True Fiber Distributed Data Interface (FDDI) is used only in a limited number of places, for medical imaging applications. While there are many fiber segments in its backbone network, M. D. Anderson only runs Ethernet protocols over them. At the end of 1994, a consortium of hospitals in the Texas Medical Center installed a true FDDI system, linking major buildings on its campus, including M. D. Anderson. This FDDI network will eventually include ATM switching. While M. D. Anderson is examining the use of fast Ethernet and FDDI, it hopes to do this on only a limited basis, leapfrogging instead to ATM.

With its current network, M. D. Anderson has achieved "universal access" – a single PC workstation on a desktop, using multiple network connections, gateways, and routers, can access any host. However, this connectivity really only addresses layers 1 through 3 of the OSI model. M. D. Anderson is still struggling with issues at layers 4 through 7.

Figure 9.1 — M. D. Anderson campus.



Managed care is beginning to exert an influence on M. D. Anderson, causing it to integrate activities as well as data. Several managed care contracts were signed in 1994. This new way of doing business requires that more people be able to access more systems, and that the data within these systems be integratable for reporting and analysis. M. D. Anderson has laid the infrastructure for this eventual integration through the installation of its backbone network. However, the middleware issues at layers 4 through 7 of the OSI model have not been resolved whereby any workstation with the proper authorization could access any data source and retrieve the desired data transparently.

9.1.2 Networks and Applications

The hospital is located in the Texas Medical Center. The two main structures are the administration building, on the south side of Holcombe Boulevard, and the main hospital on the north side of the street (Figure 9.1). A third building currently under construction will provide an additional 1.2 million square feet of facilities for patient care and research. The administration building houses an IBM 3090, which supports inpatient care and the clinics in the hospital. It is accessed by both 3270 terminals and PCs using emulation software. These devices do not access the

mainframe through the backbone network. Rather, they are connected to controllers that are directly linked by fiber cable to the front-end processors on the 3090. The administration building also houses a series of DEC VAX machines for pharmacy and diagnostic imaging. These hosts are also accessed from terminals and PCs located in the hospital. These devices, however, do use the backbone network. They are attached to a series of local area networks (LANs) in the hospital building and connected to the backbone through fiber cable and bridges.

The connectivity requirements for M. D. Anderson are staggering. Beyond the need to link all the clinical and administrative departments internally, the Information Service department has to provide communications with several dozen external entities including the University of Texas Medical School, the Texas Medical Center, University of Texas in Austin, the Harris County Hospital System, the state government in Austin, the School of Public Health, other universities, warehouses, insurers, etc. This is illustrated in Figure 9.2.

These connectivity requirements are handled with a complex of networks, consisting of a variety of cabling schemes, protocols, and network devices. This complex uses a combination of Ethernet over fiber cable, true FDDI, hard wiring, leased lines (two T-1 lines, with two more anticipated in 1995), switched lines, microwave, and Ethernet over coaxial cable. The protocols include both baseband and broadband Ethernet, FDDI, and X.25. At a higher level, the hospital uses TCP/IP, IPX, SNA, DECnet, and Appletalk. M. D. Anderson has a CAD system to maintain the documentation of its configuration. The network drawings are at several levels of detail, ranging from a high level overview down to detailed diagrams of wiring at each floor, which would include wiring closets, backbone, hubs, concentrators, and cable runs.

In order to obtain a better understanding of M. D. Anderson's connectivity requirements, we need to go to layer 7 and examine the applications, in particular the hospital hosts and LANs. This is shown in Figure 9.3.

The administration building houses most of the host computers — six DEC VAXs and an IBM 3090. One VAX is dedicated to diagnostic imaging; another is used for practice plans; a third for patient protocols; a fourth for statistical analysis; a fifth runs radiotherapy applications; and the sixth is used for pharmacy applications. The IBM 3090 runs patient registration, billing, orders, and accounting applications. Patient caregivers located in the hospital building connect to the 3090 using a combi-

nation of PCs and 3270 devices. The 3270 devices are hardwired to the 3090. They are SNA components, not really a part of the Ethernet backbone. The hospital patient monitors are not connected to any portion of the network. However, in planning for the new building, it is anticipated that the cabling infrastructure will accommodate such a connection.

Figure 9.2 — M. D. Anderson connectivity requirements.

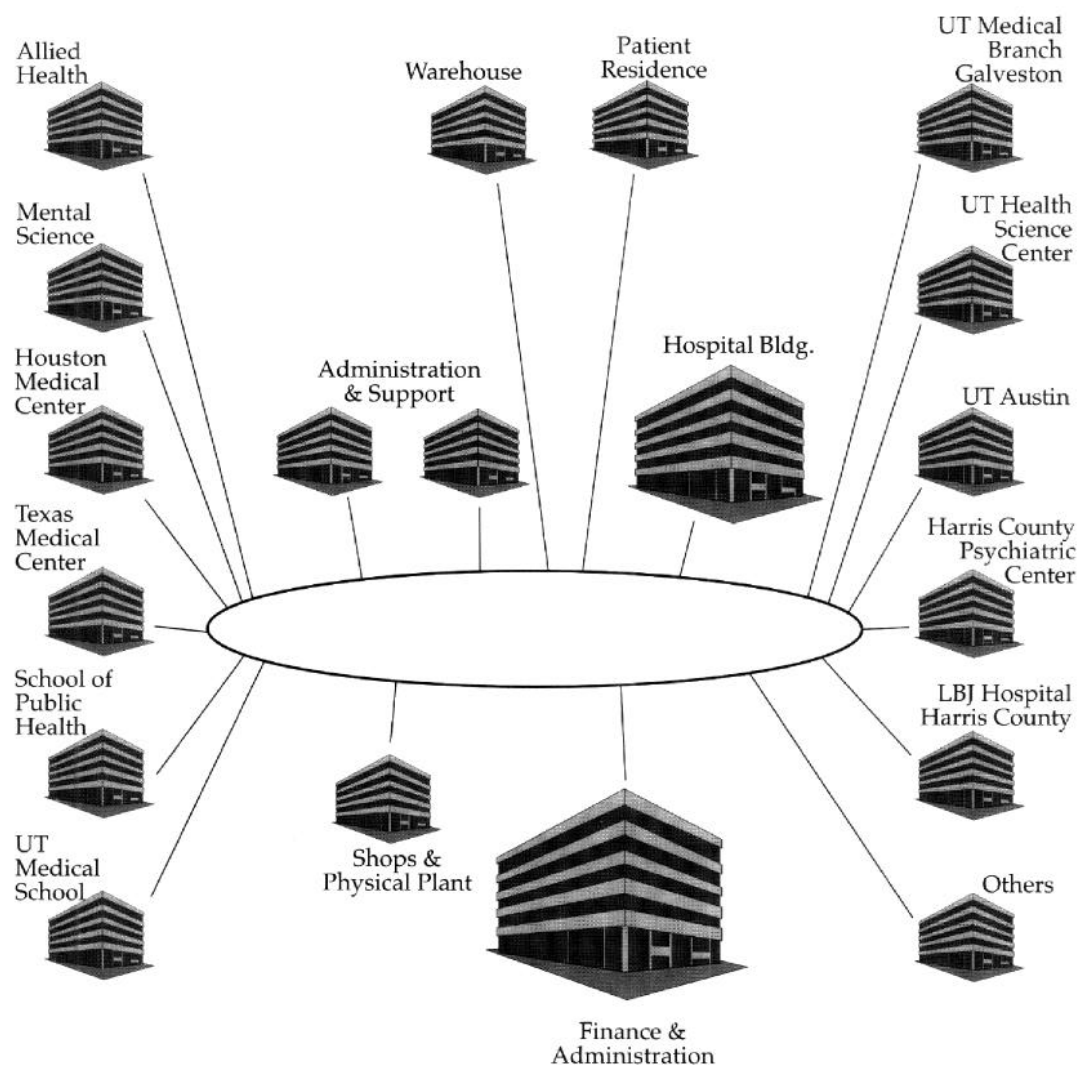
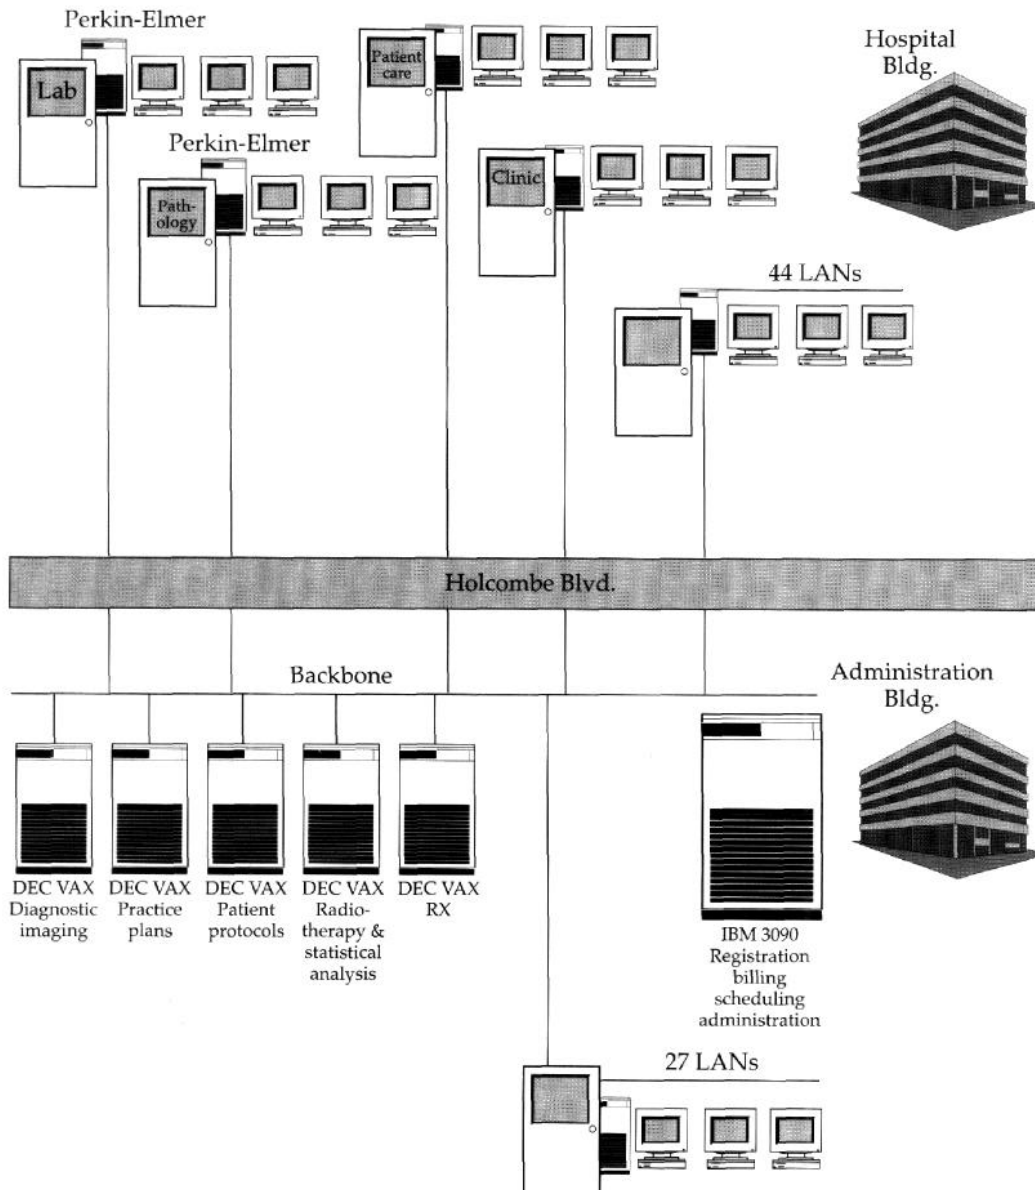


Figure 9.3 — Application level of M. D. Anderson's network.



In addition, there are 44 Novell LANs in the hospital, running a variety of departmental applications like risk management, surgery scheduling, environmental and safety, etc. There are two old Perkin-Elmer hosts in the hospital, for lab and pathology. They have terminals directly connected to them through RS232. The hosts are not located on the backbone, but can be accessed by selected 3270 devices, connected through a custom gateway built by M. D. Anderson several years ago. The backbone also has terminal servers for dumb terminals to access the VAX machines. This is done using the local area transport (LAT) protocol.

The current political and economic environment, which is fueling the growth of HMOs, and provider partnerships and consolidation, has had a significant impact on M. D. Anderson. It is actively seeking to sign HMO agreements with insurers in the Houston area and beyond. The operation of an HMO requires the integration of clinical, administrative, and financial data in a new way. This, in turn, requires that databases be integrated and that workstations have access to them. Currently, M. D. Anderson has no universal clinical workstation which could access every host database. M. D. Anderson's strategy for data access is to consolidate all relevant data to a large DB2 database on the IBM 3090. Workstations would have to access this data – both Macintosh and PC (see Figure 9.4).

In addition to the obvious issues of security and performance, this client server architecture raises the thorny question of middleware. While M. D. Anderson has resolved the basic connectivity issues up through layer 3 of the OSI model, layers 4, 5, 6, and 7 require much customization. The full complexity of this middleware is shown in Figure 9.5. PC workstations run query software, which calls programs for database access, which, in turn, make calls to Novell Netware. These requests for data are transmitted to a gateway PC on the network. This gateway machine runs OS/2, since DOS is not multitasking. This gateway must translate between the remote procedure call (RPC) used by Novell and by LU6.2 on the IBM side. It does this and sends the request to the mainframe over a token ring network to which the mainframe is connected with a 3174. The data query language is SQL. The queries are submitted to a DB2 database on the mainframe, and the data is returned to the requesting PC.

Figure 9.4 — PC clients and mainframe server.


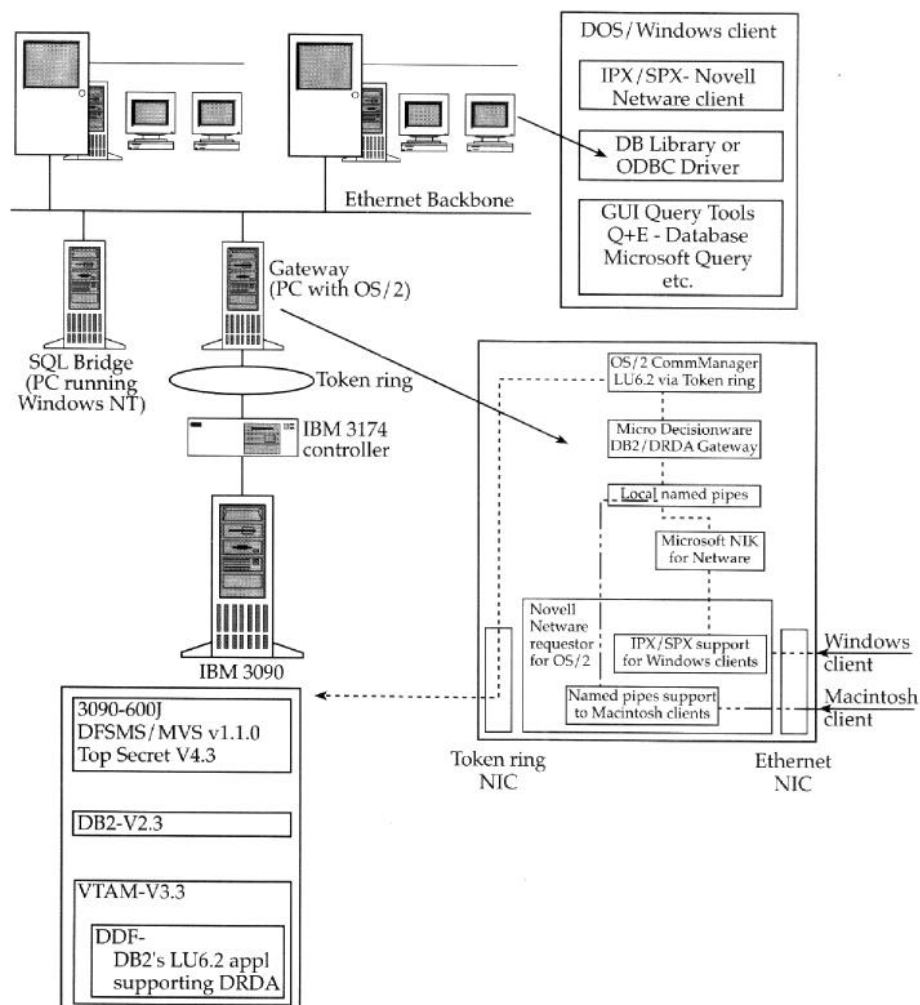


Figure 9.5 — Client/server middleware.



9.1.3 Budget and Plans

The 1994 budget for networking includes \$350,000 for salaries and \$250,000 for equipment. This does not include depreciation, which would amount to an additional \$600,000. While it was cheaper, in terms of people and equipment, to manage the previous monolithic SNA network, it could not accommodate the connectivity needs of a computing environment as diverse as M. D. Anderson's.

The Information Service Department does performance planning with end users once a year. For the SNA components of the network, it commits to a response time of less than 3 seconds for 90% of all transactions. In order to maintain this commitment, Information Systems has refrained from the practice of tunneling, the routing of SNA transactions within TCP/IP packets over the Ethernet backbone. They are reluctant to make any commitments with regard to performance on the Ethernet backbone since the software tools to manage its complexity are not sufficiently robust. This leads to certain conflicts with end users who want to run more and more of their critical applications over the backbone. Management is resisting such requests until the tools become available. In the long run, M. D. Anderson plans to move to TCP/IP universally.

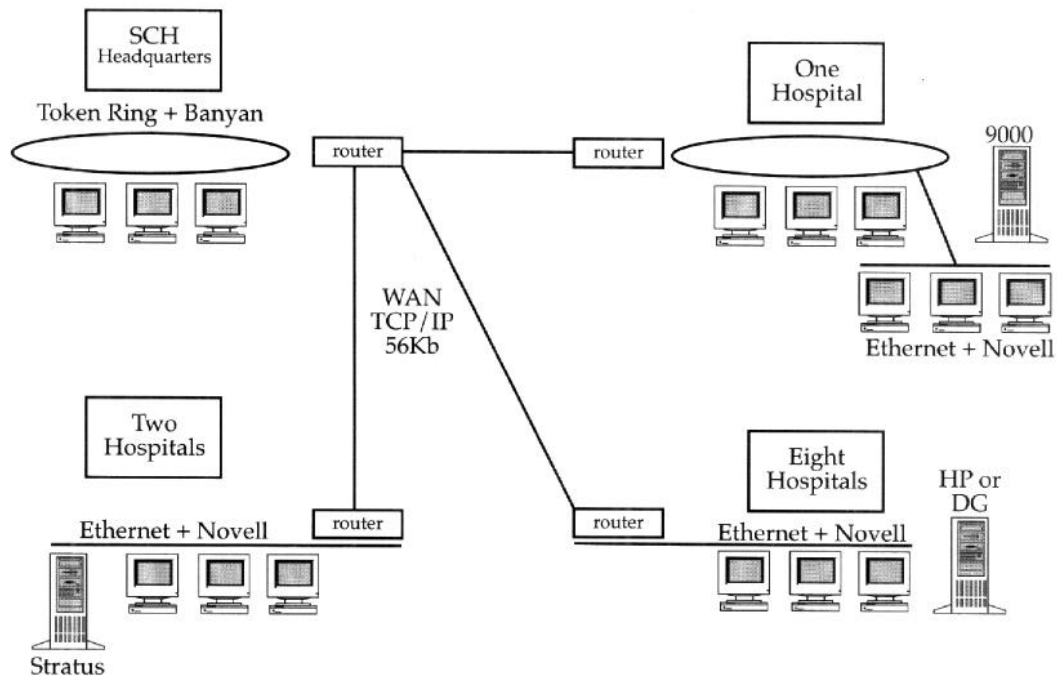
M. D. Anderson does usage accounting for CICS users of the SNA segments, but not for the Ethernet components — there is too much data and too few tools. It has a variety of software tools for the other aspects of network management: DEC software for analysis of the bridges, IBM's Netview, LattisNet from Synoptics (based on OpenView), and Sniffer. Sniffer is used to set traps for broadcast storms. The protocol for management of network devices is simple network management protocol (SNMP). As of 1991, all network devices (hubs, bridges, routers, etc.) that are acquired must contain SNMP agents.

9.2 ***Sisters of Charity***

Sisters of Charity (SCH) of the Incarnate Word Health Care System is a Catholic non-profit healthcare chain headquartered in Houston, Texas, which owns and operates 11 acute care hospitals and four long-term care centers located in Texas, Louisiana, Arkansas, California, Utah, and Ireland, with a total of 5,320 beds. In 1993, SCH had 130,248 admissions, 1,344,188 outpatient visits, and 341,523 emergency visits.

Until 1990, SCH used an IBM mainframe to support hospital operations. A centralized data processing facility in Houston, Texas provided computing services to each hospital in the chain. In the late 1980s, SCH decided to decentralize computing in order to provide greater autonomy to each hospital. This increased autonomy was a response to the changing structure of the healthcare industry, with increased consolidation of facilities, growing pressure to contain costs, prospects for national reform of the system, and pressure to incorporate more managed care principles into hospital operations.

Figure 9.6 — Sisters of Charity WAN.

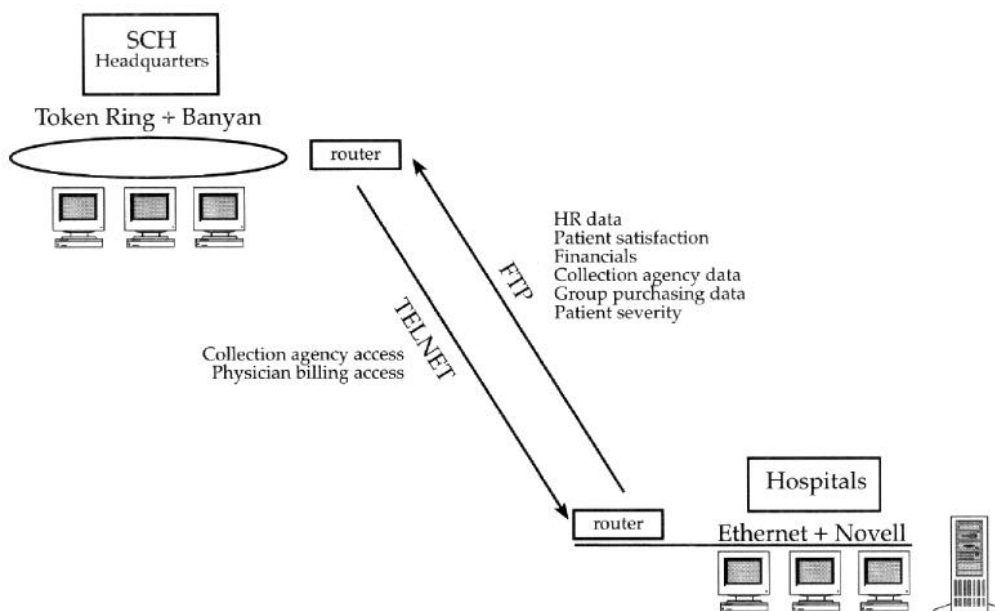


In this process of selecting a new hospital information system (HIS), eight of the 11 hospitals selected a system from HBO, which provides a suite of applications divided into financial (registration, billing, collections) and clinical (orders, pharmacy, lab, nursing plans, orders, etc.). Three of the hospitals run HBO on HP9000 machines using the UNIX operating system, while five run the system on Data General Avion machines, also using UNIX. In addition, one hospital uses a package from TDS on an IBM mainframe, and two others use MSA on a Stratus computer.

The SCH chain created a wide area network (WAN) to connect its hospitals with corporate headquarters (see Figure 9.6). At the facility level, SCH corporate headquarters performs networking differently than its hospitals. It runs the Banyan network operating system (NOS) over a token ring, while the hospitals use Novell over Ethernet. Cisco routers are used to tie the WAN together. A T-1 line ties the headquarters building to Southwestern Bell; from there, a 56 Kb line goes to each hospital. The WAN protocols are TCP/IP. Corporate headquarters requires that the hospitals provide a variety of data for consolidation and comparison: human resource data, financials, patient satisfaction surveys,

collection agency data, group purchasing data, and patient severity data. There is a corporate-wide policy governing this complex data transfer which prohibits any user from placing data on another system. If someone requires data from another system, the users of this other system extract the data and leave a file on disk. It is then the responsibility of the first user to get this data through file transfer protocol (FTP) and Telnet. For example, in order to obtain patient satisfaction data, SCH headquarters establishes a connection with the hospital's system and performs a file transfer with FTP. In order to access collection data and physician billing data from the hospital machine, SCH headquarters logs on to the hospital applications with Telnet (see Figure 9.7).

Figure 9.7 — Interaction between SCH headquarters and hospitals.



9.2.1 Corporate Headquarters

The headquarters building has structured wiring, installed when all applications were moved off the mainframe. The logical token ring is physically implemented with fiber cable in the vertical risers and

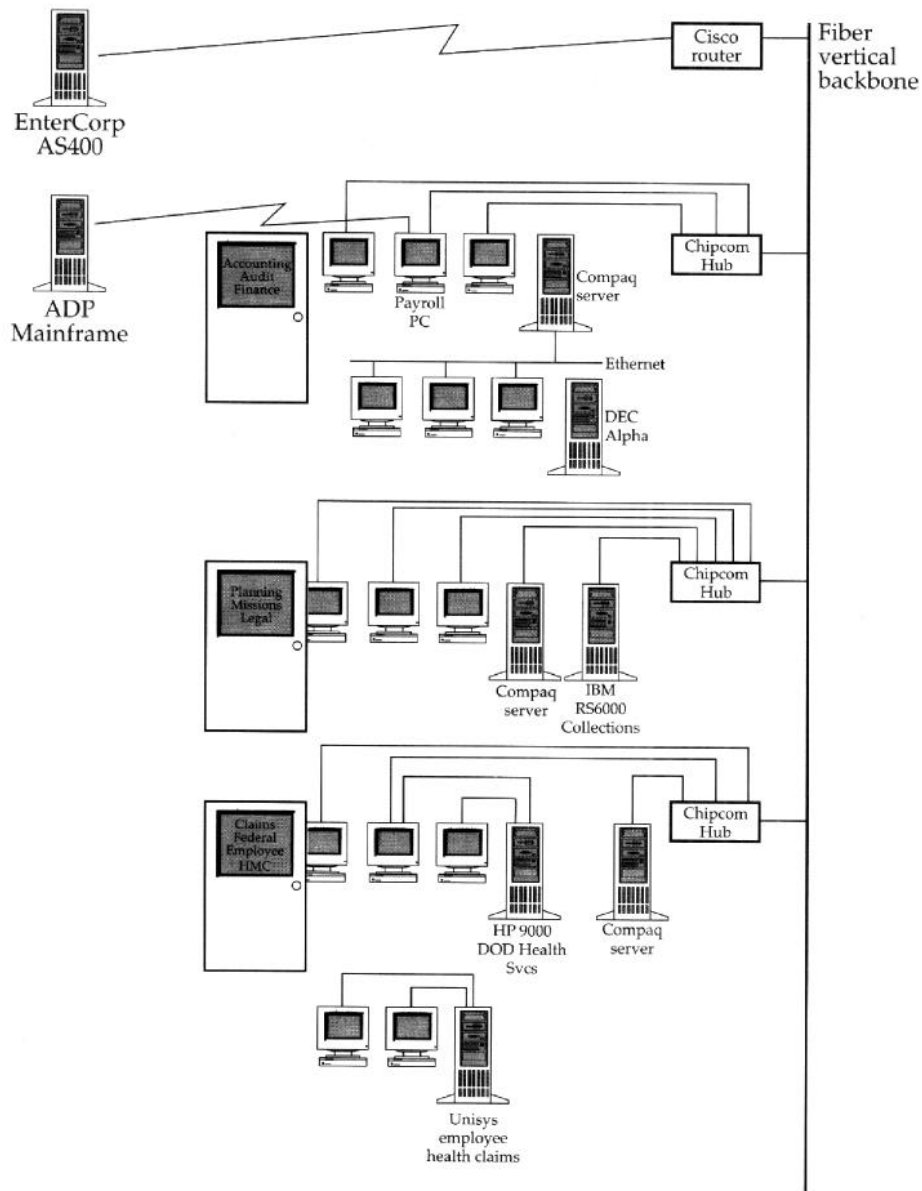
twisted-pair cable in the horizontal runs to the desktop. The twisted-pair cable is of several types (IBM Type 1, 2, and 3). All horizontal runs go back to wiring closets containing Chipcom hubs. There is no centralized data center. Instead, seven PC servers (Compaq and IBM) are attached to the token ring and located in the functional areas they serve.

In addition to the servers, there are several other platforms and connections. A DEC Alpha machine is used for accounting applications; its terminals are connected through a small Ethernet segment, which is bridged to the token ring backbone through a Compaq server. Several OS/2 machines are used for finance and decision support. A RS6000 machine running AIX is used by the collections department. Physician billing is done by a third party on a AS4000, connected through the Cisco router. A Unisys machine is used to process employee health claims. It is a standalone machine using directly connected terminals. Finally, there is an HP9000, used for a Department of Defense HMO run by SCH. This system is also standalone with PC workstations directly connected running terminal emulation software. These workstations also have a network card to access the backbone. The entire corporate headquarters network is shown in Figure 9.8.

The SCH corporate token ring network is managed with SunNet Manager (including usage accounting, failure detection, load balancing, configuration management, and security) on a SUN SPARC workstation. Cisco Works is used for router management. All hubs and routers have software with SNMP agents. In 1994, network costs were \$4,700 per month. The previous SNA network was more expensive (about \$19,000 per month); on the other hand, it had extra circuits and equipment to accommodate disaster recovery and backup.

The SCH chain plans to install Groupwise (formerly WordPerfect Office) for corporate-wide e-mail and calendaring, while continuing to use FTP for file transfer. It is also examining the feasibility of installing a PC-based videoconferencing system to connect corporate headquarters with the hospitals using T-1 lines. In response to the changing healthcare environment, SCH has entered into a partnership with another Houston chain, Memorial Health Care System, to provide a healthcare network for southeast Texas. At the end of 1994, this venture was in the process of formation, and there were no new requirements for network services, but eventually this partnership will have a dramatic impact on the nature and structure of networking at SCH.

Figure 9.8 — SCH corporate network.



9.2.2 St. Mary's Hospital

St. Mary's Hospital, a 322-bed adult critical care hospital located in Galveston, Texas, is part of the SCH chain. Until 1992, St. Mary's utilized the mainframe computer located in Houston for accounting, registration, and patient care. When the decision to decentralize computing was made, St. Mary's, like most of the SCH hospitals, selected the HBO clinical and financial applications to be run locally. For the hardware platform, St. Mary's chose the 9000/867. For its hospital network, St. Mary's selected Novell running over Ethernet. In preparation for the HBO installation, St. Mary's rewired its entire facility, which consists of a single contiguous building with four towers. Fiber cable was run vertically in each tower, and horizontally to connect the towers. Within each tower, Level 5 cabling was installed between the wiring closets and desktops. This is shown in Figure 9.9.

The fiber cable runs through wiring closets, where Hughes hubs are located. The main hub, which is located in the Data Center, is a Hughes 1000 MDF (main distribution facility). The wiring closets on each floor contain Hughes 1100 IDF (intermediate distribution facility) hubs, and Hughes terminal servers for asynchronous devices such as card readers, lab instruments, and asynchronous printers. In order to improve network reliability, St. Mary's installed redundant fiber cable connecting the MDF with the IDF units. The Data Center houses most of the host computers and LAN servers. This configuration is shown in Figure 9.10.

Figure 9.9 — St. Mary's cabling scheme.





Figure 9.10 — Application level view of St. Mary's network.



Network management (failure detection, configuration, and load balancing) is done through a Hughes MONET workstation. However, in the future, St. Mary's will use SunNet Manager which will be bundled with the network management workstation. The network and applications cost \$4 million to install. Costs for 1994 totaled \$1.4 million.

The large administrative and clinical applications are from the HBO and run on the HP9000 computers. All these applications share the same database. Other applications, like surgery and blood bank, are run on PC servers. St. Mary's does not use HL7 to synchronize the HBO databases with the surgery database on the surgery server since the existing interface does not support HL7. Instead, dual data entry is performed. A DEC 4000 is used to run HBO's Trendstar decision-support package. They run the UNIX operating system, with TCP/IP as the network protocol. The other servers (payroll, office automation, and surgery) run Novell with IPX on PC workstations. Users access the applications on computers through TCP/IP and terminal emulation; for office automation and PC applications they must use IPX. Some workstations, therefore, must load two protocol stacks simultaneously since some users need access to multiple systems. St. Mary's does have universal access. All hosts can be accessed through a single workstation. Security is handled by SaberLan. Although all users currently have PCs, consideration is being given to the use of dumb terminals for "heads down" users – employees who access a single application exclusively in the performance of their job.

St Mary's ICU patient monitors are not connected to any of the information systems. Currently, there are no imaging applications to support radiology services. While the hospital is located within a major medical center (the University of Texas Medical Branch at Galveston), it does not have any connections with other networks in this complex. Physicians with admitting privileges at St. Mary's can dial in from their office to view patient data.

Eventually the hospital would like to move to ATM switches for increased bandwidth, ease of configuration, and load balancing.

10.0 ABBREVIATIONS

ACR-NEMA	American College of Radiology - National Electrical Manufacturers Association	GUI	graphical user interface
ANSI	American National Standards Institute	HAP	healthcare automation protocol
API	application program interface	HBOC	HBO & Company
ARPANET	Advanced Research Projects Agency Network	HL7	health level seven
ASTM	American Standard for Testing and Materials	HIS	hospital information system
ATM	asynchronous transfer method	HMO	health maintenance organization
AUI	attached unit interface	IBM	International Business Machine
AWG	American Wire Gauge	ID	identification
BCC	bedside communications controller	IDF	intermediate distribution facility
BBS	bulletin board system	IEC	International Electrotechnical Commission
BSI	British Standard Institute	IEEE	Institute of Electrical and Electronic Engineers
CAD	computer-aided design	IP	internet protocol
CICS	Customer Information Control System	IPX	internet packet exchange
CIS	clinical information system	IS	intermediate system
CCITT	Consultative Committee for International Telegraph and Telephone	ISA	industry standard architecture
CHIN	community healthcare information network	ISDN	integrated services digital network
CLNP	connectionless network protocol	ISO	International Standardization Organization
CMIP	common management information protocol	ISP	international standards profile
CMIS	common management information service	ITU	International Telecommunications Union
COAX	coaxial cable	ITU-TSS	International Telecommunications Union - Telecommunications Standardization Sector
COM	common object model	Kb	kilobyte
COS	Corporation for Open Systems	LAN	local area network
CMS	cable management system	LAT	local area transport
CPU	central processing unit	LCC	logical link control
CSMA/CD	carrier sense multiple access with collision detection	LED	light emitting diode
DBMS	database management system	LU	logical units
DCC	device communications controller	MAC	media access control
DCE	distributed computing environment	MAN	metropolitan area network
DEC	Digital Equipment Corporation	MAP	manufacturing automation protocol
DDE	dynamic data exchange	MAU	multistation access unit
DME	distributed management environment	Mbps	megabits per second
DOS	disk operating system	MCA	microchannel bus
EDI	electronic data interchange	MDF	main distribution facility
EIA	Electronic Industries Association	MEDIX	medical data interchange
EISA	enhanced industry standard architecture	MIB	medical information bus
e-mail	electronic mail	MONET	managing open networks
EMI	electromagnetic interference	NFS	network file system
FDDI	fiber distributed data interface	NIC	network interface card
FTAM	file transfer access and management	NOS	network operating system
FTP	file transfer protocol	ODBC	open database connection
		OLE	Microsoft's object linking and embedding
		OS	operating system
		OSF	Open Systems Foundation

OSI open systems interconnection
OSPF open shortest path first
PC personal computer
PCI peripheral component interconnect
PCMCIA Peripheral Computer Memory Card International Association
PDA personal data assistant
PO purchase order
PVC polyvinyl chloride
RAID redundant arrays of inexpensive disks
RAM random access memory
RDA remote database access
RFI radio-frequency interference
RIP routing information protocol
RISC reduced instruction set computing
ROSE remote operations service element
RPC remote procedure call
SAA systems application architecture
SCH Sisters of Charity Hospitals
SCSI small computer system interface
SDLC synchronous data link control
SMTP simple mail transfer protocol
SNA systems network architecture
SNMP simple network management protocol
SPARC scalable processor architecture
SPX sequenced packet exchange
SQL structured query language
STP shielded twisted-pair
TCP transmission control protocol
TCP/IP transmission control protocol/internet protocol
TOP technical office protocol
UPS uninterruptible power supply
UT University of Texas
UTP unshielded twisted-pair
VAC value added carrier
VESA Video Electronics Standards Association
VFS virtual file system
VIM vendor independant messaging
VINES vitrual networking software
VMS virtual memory system
WAN wide area network
XNS Xerox network services

11.0 BIBLIOGRAPHY

The following bibliography lists citations pertinent to networks.

Corrinag PH, Guy A. *Building Local Area Networks*. New York, NY: M & T Books; 1992.

Derfler FJ. *Guide to Connectivity*. Emeryville, CA: Ziff Davis Press; 1992.

Held G. *Internetworking LANs and WANs*. New York, NY: John Wiley and Sons; 1993.

Heywood D, Jerney J, Johnston J, et al. *Enterprise Series: LAN Connectivity*. Indianapolis, IN: New Riders Publishing; 1992.

Hunter P. *Networking Operating Systems Making the Right Choices*. Redding, MA: Addison-Wesley; 1994.

Madison T. *Local Area Networks New Technologies, Emerging Standards*. New York, NY: John Wiley and Sons; 1988.

Naugle M. *Network Protocol Handbook*. New York, NY: McGraw Hill, Inc.; 1994.

Rosch WL. *The Winn L. Rosch Hardware Bible*. New York, NY: Brady Publishing; 1992.

Rose MT. *The Open Book, A Practical Perspective on OSI*. Englewood Cliffs, NJ: Prentice Hall; 1990.

Sheldon T. *LAN Times Encyclopedia of Networking*. Berkeley, CA: Osborne McGraw-Hill; 1994.

Taylor A. A serial interface design to integrate bedside devices into patient monitoring systems. *J Clin Eng*. 1992; 17:325-329.

12.0 GLOSSARY

ACR/NEMA — Its goal is to develop a standard for connecting digital imaging devices.

ARCnet — Like token ring, ARCnet utilizes tokens to notify network components when they can transmit their data packets or frames. Each ARCnet station is identified by a unique address (from 1 to 255) that is usually set on the NIC by the network installer.

application layer — The application layer relies on services performed at lower levels but is the layer least involved with the underlying network hardware. Tasks performed on the application layer vary with the uses of a network, but they might include login procedures, electronic mail, terminal emulation, database management, and the operation of file servers and print servers.

application programming network — A convention that allows heterogeneous peer-to-peer networking over SNA networks.

bridges — Connects similar or identical local area networks. These connections occur at the media access control sublayer of the OSI model.

bus topology — Connects network components in a line, one after the other. Each network component taps into the cable as it snakes its way along from one component to another.

carrier sense multiple access with collision detection — A network protocol for handling situations in which two or more nodes transmit at the same time, thus causing a collision. To avoid another collision, both then wait for differing random amounts of time before attempting to transmit again.

client/server — Characterized by microcomputers operating as intelligent client workstations connected to a computer operating as a server.

Coax cable — Is constructed with a core copper wire covered by plastic insulation which is surrounded by a metallic foil or woven mesh copper shield, followed by yet another layer of insulation as an outer jacket.

concentrator — A communications device that combines signals from multiple sources, such as terminals on a network, into one or more signals before sending them to their destination.

data link layer — Is one level above the physical layer; it is involved both in packaging and addressing information and in controlling the flow of separate transmissions over communication lines.

electronic data interchange — The ability to transfer information such as orders and invoices from one computer to another over a communications network. The goal of EDI is to eliminate the redundant paperwork and delays in response time inherent in mail and delivery services.

electronic mail — Is an application that enables a user to broadcast messages with attached files to an individual or group on the network.

Ethernet — A local area network developed by Xerox in 1976, originally for linking minicomputers at the Palo Alto Research Center. A widely implemented network from which the IEEE 802.3 standard for contention networks was developed, Ethernet uses a bus topology and relies on the form of access known as CSMA/CD to regulate traffic on the main communications line.

fiber distributed data interface — A standard developed by the American National Standards Institute for high-speed fiber-optic local area networks. FDDI provides specifications for transmission rates of 100 megabits per second on networks based on the token ring standard.

file transfer — The process of moving or transmitting a file from one location to another, as between two programs or from one computer to another.

gateway — A device used to connect dissimilar networks - networks using different communications protocols - so that information can be passed from one to the other.

health level 7 — The HL7 working group purpose is to develop standards governing the exchange of key data sets among healthcare applications.

host-terminal — Is characterized by a host (mainframe, mini or micro) computer with dumb terminals attached.

hub — A simple device housed in a cabinet which connects nodes via a specific type of cabling.

integrated services digital network — A worldwide digital communications network evolving from existing telephone services. The goal of the ISDN is to replace current telephone lines, which require digital-to-analog conversions, with totally digital switching and transmission facilities capable of carrying data.

International Standards Organization — ISO is located in Geneva, Switzerland, and interacts with the standards bodies of individual countries through a hierarchical network of relationships.

Internet — In communications, a set of computer networks - possibly dissimilar - joined together by means of gateways that handle data transfer and the conversion of messages from the sending network to the protocols used by the receiving network (with packets if necessary).

Internet protocol — It contains protocols for terminals, file transfer, and session management for the Internet.

logical link control — The LLC applies to all IEEE 802 standards and covers station-to-station connections, generation of message frames, and error control.

media access control — The MAC differs from one IEEE 802 standard to another and deals with network access and collision detection.

Medical Information Bus — The IEEE established the MIB in 1984. The objective of this standards committee is to establish a uniform method of communication between medical bedside devices and patient computer systems that are utilized primarily in acute care environments.

MEDIX — The goal of MEDIX is to define standards for the interoperability of heterogeneous healthcare information systems, in contrast to HL7, which focuses on the narrower issue of interfaces.

metering software — Local area network software that monitors the use of applications by network stations and provides a report of the usage history.

middleware — Using a simplified application programming interface, the developer makes a call to the middleware, which handles all the other calls to network modules, program libraries, etc.

multistation access unit — Typically used to refer to a token ring wiring center. It functions similar to hubs and concentrators.

network interface — The physical connection between the computer and the network cabling system.

network file system — Network file system makes remote files and directories appear to be part of the local system.

network layer — The network layer is one level above the data-link layer and ensures that information arrives at its intended destination. The function of the network layer is to establish, maintain, and keep open a path for information to travel on and to make the actual route immaterial to any other layer.

network operating system — An operating system installed on a server in a local area network that coordinates the activities of providing services to the computers and other devices attached to the network.

open systems interconnection — A layered architecture that standardizes levels of service and types of interaction for computers exchanging information through a communications network. The OSI model separates computer-to-computer communications into seven layers, or levels, each building upon the standards contained in the levels below it.

peer-to-peer — Characterized by two or more microcomputers connected together and running network software.

personal data assistant — A combination of wireless phone and computer that can fit in a person's hand.

RAID — Redundant arrays of inexpensive disks that provide some hard disk subsystem redundancy.

remote database access — A standard for the opening and closing of a remote database.

remote procedure call — A RPC makes it possible to execute a program that is located on another computer. This remote computer stores and runs the program. A RPC looks like a call to a subroutine on the local computer; however, the subroutine is located on the remote computer.

repeater — A device used on communication circuits that decreases distortion by amplifying or regenerating a signal so that it can be transmitted onward in its original strength and form. On a network, a repeater connects two networks or two network segments at the physical layer of the OSI model and regenerates the signal.

ring — Connects each network component to both the components before and after it on the ring. Signals are typically passed along from one component to the next in one direction along the ring rather than broadcast.

routers — An intermediary device on a communications network that expedites message delivery.

routing information protocol — RIP is a routing algorithm based on counting hops (intermediate nodes between sender and addressee). The sending node decides whether the destination node is on the same subnet-work. If not, the sending node forwards the message to a router, which, in turn, determines the best route.

sequenced packet exchange — A protocol enabling program-to-program communication, error correction, check pointing, and flow control.

server — On a local area network, a computer running administrative software that controls access to all or part of the network and its resources.

session layer — The session layer handles the details that must be agreed upon and followed by two devices exchanging information; it coordinates and regulates the transfer of information and maintains the session for as long as needed.

simple mail transfer protocol — SMTP is a store-and-forward model, that is, when a mail message is transmitted by the sender, it is relayed to an intermediate computer, where it is stored until it can be forwarded to the recipient's computer. When the message arrives at the recipient's computer, it is placed in a queue, and later moved to the recipient's mailbox storage area.

simple network management protocol — The SNMP was developed for the TCP/IP to managed the multivendor networks. It is based upon the concept of agents and management information bus.

star — Connects each network component to a wiring center to propagate signals between components.

systems network architecture — A widely used communications framework developed by IBM to define network functions and establish standards for enabling its different models of computers to exchange and process data.

thinnet — Thin Ethernet cabling is 5 millimeters in diameter and can connect network stations over a distance of 300 meters.

token ring — A local area network formed in a ring (closed loop) topology that uses token passing as a means of regulating traffic on the line. Token ring networks are defined in the IEEE 802.5 standard.

transport layer — The transport layer is one level above the network layer and is responsible for both quality of service and accurate delivery of information. Among the tasks performed on this layer are error detection and correction.

translation — The frames of one network are converted into the frames of another network, through the use of a multiprotocol router.

transmission control protocol/internet protocol — A software protocol developed by the Department of Defense for communications between computers.

tunneling — With tunneling there is no translation. The frame on one network type is encased within layer 3 frame of another network.

UNIX — A multiuser, multitasking operating system for use on minicomputers.

unshielded twisted-pair — Cable that consists only of one or more pairs of cable twisted together within an unshielded insulated jacket.

virus — A program that infects computer files by inserting in those files copies of itself. Viruses often have damaging side effects, sometimes intentionally, sometimes not.

wide area networks — Link multiple local area networks, thereby giving users access to information and resources across the combined system.

INDEX

A

American College of Radiology - National Electrical
Manufacturers Association 30, 70, 75
American National Standards Institute 8, 32, 70, 71
American Standard for Testing and Materials 30, 70, 74
ARCnet 42-43, 46, 60, 66
asynchronous transfer method 17, 68, 69, 77, 92
attached unit interface 15, 38

B

Banyan Systems 47-48, 86
bridges 7, 37, 59, 60, 62, 68, 80, 85, 89
bus 14, 30, 33-35, 38, 39, 55

C

carrier sense multiple access with collision detection 38
clinical information system 13, 75
coax cable 12, 14, 31, 38, 43, 66, 82
common management information protocol 63
Community Healthcare Information Network 68, 69-70
complex document exchange 23, 26
concentrators 36, 37, 63, 82
Consultative Committee for International Telegraph and
Telephone 8, 9, 17
Corporation for Open Systems 8, 9

D

database management system 28
data link 15-16, 59
DEC 4, 7, 19, 48, 49, 65, 77-88
distributed management environment 65

E

electronic data interchange 25, 28-30, 71
Electronic Industries Association 8
electronic mail 2, 24-27, 48, 50, 60, 69
enhanced industry standard architecture 30
Ethernet 4, 6, 9, 12, 15, 16, 18, 37-38, 42, 43, 46, 48,
59, 66, 68, 76-90

F

fault-tolerant 58, 60
fiber distributed data interface 7, 16, 32, 52, 77, 79
fiber-optic cable 32, 51
file exchange 27, 70
file transfer access and management 27
file transfer protocol 27, 87, 88

G

gateways 25-26, 35, 60, 62, 71, 77, 79, 84, 86

H

HL7, health level seven 23, 30, 70-72, 76, 92
health maintenance organization 50, 78, 84, 91
hospital information system 86
hubs 7, 15, 36-37, 40, 62, 77, 82, 85, 90, 92

I

industry standard architecture 30
Institute of Electrical and Electronic Engineers 8, 9, 16,
37-42, 72, 77
integrated services digital network 51, 68
intermediate distribution facility 90
intermediate system 17
International Electrotechnical Commission 8
International Standardization Organization 7-10, 28
international standards profile 75
International Telecommunications Union 8, 9
internet protocol 3, 16, 17, 18, 19, 21, 25, 26, 27,
63, 66, 67, 75, 77, 82, 92
internet packet exchange 11, 15, 18-19, 66, 75, 77, 92

L

local area network 2-7, 12, 15, 18, 21, 24, 31, 32, 33, 37,
41, 43, 48, 49, 51, 52, 59, 61, 62, 68, 70, 71
local area transport 16, 19, 82

M

main distribution facility 90
managing open networks 92
manufacturing automation protocol 30, 74
media access control 15, 59
medical data interchange 30, 70, 72-74
medical information bus 13, 30, 70, 75
metropolitan area network 51
microchannel bus 30
multistation access unit 15, 36, 37, 41-42

N

naming 9, 10, 24-25, 26, 27
network file system 27
network interface card 30-31, 43, 62
network operating system 7, 22, 46, 47, 48, 56, 57, 86
Netview 65, 85
Novell 4, 16, 17-19, 21, 26, 46, 52, 77, 82, 86, 90

O

open database connection 28
open shortest path first 18
Open Software Foundation 8, 9, 25, 65, 79
open systems interconnection 7, 8, 10, 12, 14-30, 59, 61,
62, 71, 72, 73-77, 82
open view 65

P

performance technology 49
peripheral component interconnect 30, 52
Peripheral Computer Memory Card International
Association 30
personal data assistant 6

R

reduced instruction set computing 47
redundant arrays of inexpensive disks 58-59
remote database access 28
remote operations service element 27, 28, 72
remote procedure call 25, 27, 68, 82
repeater 36, 37
resource sharing 49-50
ring 14, 32, 34, 36, 41
router 7, 12, 17, 19, 34, 37, 60, 62, 63, 68, 76, 77, 85, 88
routing information protocol 18

S

scalable processor architecture 90
security 9, 23, 25, 48, 49, 58, 62, 65, 68, 85, 90
server 7, 12, 18, 26, 31, 40, 44-50, 52-55, 67, 84
sequenced packet exchange 19, 21
shielded twisted-pair 32, 41, 42
simple mail transfer protocol 25-26
simple network management protocol 7, 62-63, 85, 88
Sisters of Charity Hospitals 85-92
star 15, 34-36, 40-42
structured query language 28, 82
synchronous data link control 12
systems application architecture 77
systems network architecture 15-19, 65, 77, 84, 88

T

technical office protocol 30, 74
token ring 15, 16, 41-43, 46, 60, 66, 68, 87-88
topology 15, 33-35, 38-42, 49, 64
transmission control protocol 3, 16-19, 21, 25, 26, 27,
62, 63, 74, 77, 82, 85
tunneling 17, 19-20, 85

U

UNIX 22, 49, 65, 86
unshielded twisted-pair 32, 40, 42

V

Video Electronics Standards Association 30
virtual circuits 17
viruses 67

W

wide area network 2-3, 19, 21, 24, 41, 50, 59-61, 68-71, 86

X

Xerox network services 11, 19, 66



Spacelabs Medical, Inc.
15220 NE 40th Street, P.O. Box 97013
Redmond, WA 98073-9713
(425) 882-3700

ISBN 0-9627449-9-9
P/N 061-0417-00